Bites of Tomorrow: Personalized Recommendations for a Healthier and Greener Plate

Jiazheng Jing*, Yinan Zhang*, Chunyan Miao†

College of Computing and Data Science (CCDS), Nanyang Technological University (NTU), Singapore
jiazheng001@e.ntu.edu.sg, {yinan.zhang, ascymiao}@ntu.edu.sg

Abstract

The recent emergence of extreme climate events has significantly raised awareness about sustainable living. In addition to developing energy-saving materials and technologies, existing research mainly relies on traditional methods that encourage behavioral shifts towards sustainability, which can be overly demanding or only passively engaging. In this work, we propose to employ recommendation systems to actively nudge users toward more sustainable choices. We introduce *Green Recommender Aligned with Personalized Eating* (*GRAPE*), which is designed to prioritize and recommend sustainable food options that align with users' evolving preferences. We also design two innovative *Green Loss* functions that cater to green indicators with either uniform or differentiated priorities, thereby enhancing adaptability across a range of scenarios. Extensive experiments on a real-world dataset demonstrate the effectiveness of our *GRAPE*.

Introduction

The recent emergence of extreme climate events, including the record-breaking heatwaves in July 2024, catastrophic floods, and increasingly frequent wildfires, indicates a troubling shift toward global instability and highlights the critical urgency of addressing climate change (Jain et al. 2022; Ngcamu 2023). Significant research efforts have been conducted to develop proactive solutions towards sustainability, including the development of renewable technologies and materials that enhance energy efficiency (Samir et al. 2022; Fang and Durable 2024). Additionally, innovations in water purification, desalination, and recycling technologies are being explored to ensure a sustainable water supply(Darvishi et al. 2023; Curto, Franzitta, and Guercio 2021).

Since human activities play an important role in global sustainability, some research works concentrate on understanding and encouraging behavioral shifts towards sustainability (Kirby and Zwickle 2021; Marcus and Roy 2019). For example, (Štofejová et al. 2023) explore the influence of factors such as environmental attitudes and lifestyle on online purchasing behavior, identifying significant impacts from an environmentally oriented lifestyle and willingness to pay for green products. Regarding behavioral nudges, most existing works focus on implementing regulations and designing poster campaigns to promote sustainable practices (Peleg Mizrachi and Tal 2022). These methods, while traditional, can sometimes be perceived as imposing or may rely on passive engagement, which might not consistently engage or motivate the target audience.

Recommendation systems have shown considerable potential to deliver effective, personalized, and timely nudges that can lead to meaningful shifts in user behavior across various domains (He, Liu, and Jung 2024). Given the vast amount of products and services available, it is often impractical for consumers to review the whole item pool. Instead, they typically make choices from a curated list of options presented to them. By learning from users' historical behaviors and ongoing feedback, recommendation systems can prioritize sustainable items that align with user preferences, thereby effectively guiding consumers towards more sustainable choices (Zhang et al. 2024a). However, most existing works simply aim to improve prediction accuracy or enrich user experience, often neglecting the potential of recommendation systems to influence user behavior towards sustainability (Zhang et al. 2024b; Lin et al. 2024).

In this work, we introduce a novel task for food recommendation systems, named *Green Food Recommendation*: to recommend foods that not only appeal to users but also actively prioritize and promote more sustainable options. We focus on the food domain for several reasons. First, food production is associated with multiple environmental issues, including water pollution, carbon emissions, and the use of arable land (Clark et al. 2022). Second, food is a fundamental part of daily life, and even minor sustainable improvements can lead to significant impacts (Poore and Nemecek 2018). Third, food sustainability encompasses a variety of criteria, including environmental impact, nutritional value, and more. By integrating multiple sustainability indicators, the recommendation approaches becomes more robust and versatile, applicable across various domains.

Several challenges need to be addressed in *Green Food Recommendation*. First, individual attitudes towards different aspects of sustainability can vary and change over time. Simply increasing the exposure to more sustainable food options may lead to user dissatisfaction. Therefore, it is crucial to recommend items that balance both evolving personal

*These two authors contributed equally to this work.
†Corresponding author.

preference and sustainability. Moreover, the multiple indicators associated with food sustainability complicate the optimization of recommendations. These indicators, which reflect various aspects of food, can vary significantly. For example, although fish is nutritious, it often scores poorly on environmental impact indicators due to the water pollution caused during its cultivation.

To tackle the aforementioned challenges, we propose a novel green food recommendation method named *Green Recommender Aligned with Personalized Eating* (*GRAPE*). To model users' varying preferences and attitude towards sustainability, we employ both self-attention and cross-attention mechanisms. To encourage more sustainable food choices, we designed two novel *Green Loss* functions that prioritizes more sustainable options. Regarding the *Green Loss* functions, one treats all sustainability indicators equally, while the other assigns them different priorities. Extensive experiments on a real-world food dataset (Zhang et al. 2024a) were conducted to evaluate the proposed *GRAPE*. Experimental results demonstrate that *GRAPE* outperforms all state-of-the-art baselines in making both accurate and sustainable recommendations.

With this paper, we make the following contributions:

- We propose a novel research task of *Green Food Recommendation* to encourage further research that not only caters to user preferences but also nudges user behaviors towards sustainability.
- To bridge the gap between users' preferences and item sustainability, we propose *Green Recommender Aligned with Personalized Eating* (*GRAPE*). In *GRAPE*, we designed two innovative *Green Loss* functions that address green indicators with either equal or differentiated priorities, enhancing adaptability in various scenarios.
- Extensive experimental results highlight the superior performance of *GRAPE* in making accurate and sustainable recommendations. Furthermore, our findings demonstrate that it is feasible to maintain recommendation accuracy while enhancing the sustainability attributes of the recommended food. We believe that our research establishes a solid groundwork for future studies in this area.

Related Work

Recent years have witnessed the worsening of environmental issues such as global warming and rising sea levels, which has brought increasing attention to sustainability issues from academia.(Jain et al. 2022; Ruggerio 2021). Many studies focus on investigating individual environmental awareness, for example, to explore the environmental awareness between people from different backgrounds (Kirby and Zwickle 2021), and investigate the causes and future trends of people's sustainability awareness (Marcus and Roy 2019). Besides, there are also researches focusing on exploring the current status and influence of specific environmental issues, for instance, climate change (Orlinska-Wozniak et al. 2021), mountain fire (Yin et al. 2024), and so on. Additionally, solutions to environmental problems are also a key focus of research across various fields, including environmentally friendly materials (Samir et al. 2022) and industrial emission reduction (Fang and Durable 2024). These studies have made significant contributions to a sustainable society. However, through exploring the importance of individual behavior in addressing environmental issues, these studies have rarely delved deeply into the possibilities of encouraging people to live greener lifestyles.

Artificial intelligence (AI) has developed rapidly in recent years, showing great potential to solve problems in various fields, including leverage AI to solve sustainability problems, called Green AI (Schwartz et al. 2020; Verdecchia, Sallou, and Cruz 2023). For example, semantic, segmentation and multimodal methods are utilized for satellite image processing (To et al. 2024) and climate forecast (Jones 2017), conserving resources while also improving task accuracy. These efforts provide efficient solutions to specific problems. However, few works consider leveraging AI's influence on people's behavior to address sustainability issues.

Recommendation systems (RS), which are widely applied in modern society, have great potential to influence individual behaviors, which means that recommendation systems could be able to encourage and nudge users to adopt a greener lifestyle (He, Liu, and Jung 2024; Forouzandeh et al. 2024; Rostami, Aliannejadi, and Oussalah 2023). Notice this, Zhang et al. (Zhang et al. 2024a) introduced a green dataset for food recommendations where each food is labeled with three sustainability indicators. Based on this dataset, CLUSSL (Zhang et al. 2024b) is proposed to improve the recommendation performances utilizing sustainability indicators. However, most of the work focuses only on improving recommendation accuracy, with very few efforts dedicated to providing greener recommendations.

Method

Green Food Recommendation Task

Given an item $i \in \mathcal{I}$, there are n sustainability indicators associated with i, labeled as $g_1^i, g_2^i, ..., g_n^i$, respectively. Given user set $\mathcal{U}$ with size $|\mathcal{U}| = v$, for user $u \in \mathcal{U}$ who have interacted with w items, we denote the interacted item sequence as $\mathbb{I}^u = [i_1, i_2, ..., i_w]$, where i_j represents the j-th interacted item in the chronologically ordered. Similarly, the sequence of the corresponding sustainability indicators for $\mathbb{I}^u$ is denoted as $\mathbb{G}_1^u = [g_1^{i_1}, g_1^{i_2}, ..., g_1^{i_w}], ..., \mathbb{G}_n^u = [g_n^{i_1}, g_n^{i_2}, ..., g_n^{i_w}]$. The goal of the *Green Food Recommendation* is to recommend items that not only satisfy the user's interests but also have high sustainability indicators.

Model Architecture

Figure 1 shows the overall structure of our proposed *GRAPE*, which consists of: Embedding Module, Sustainability Integrated Attention Module, and Prediction Module.

Embedding Module In this module, we introduce two distinct embedding functions for the interacted items and their corresponding sustainability indicators. First, we embed the sequence of items $\mathbb{I}^u$ that user u has interacted with, using a one-hot embedding layer $\mathbf{L}_I$ to obtain embeddings $E_i^0 \in \mathbb{R}^{w \times d}$. Here, d represents the embedding size.

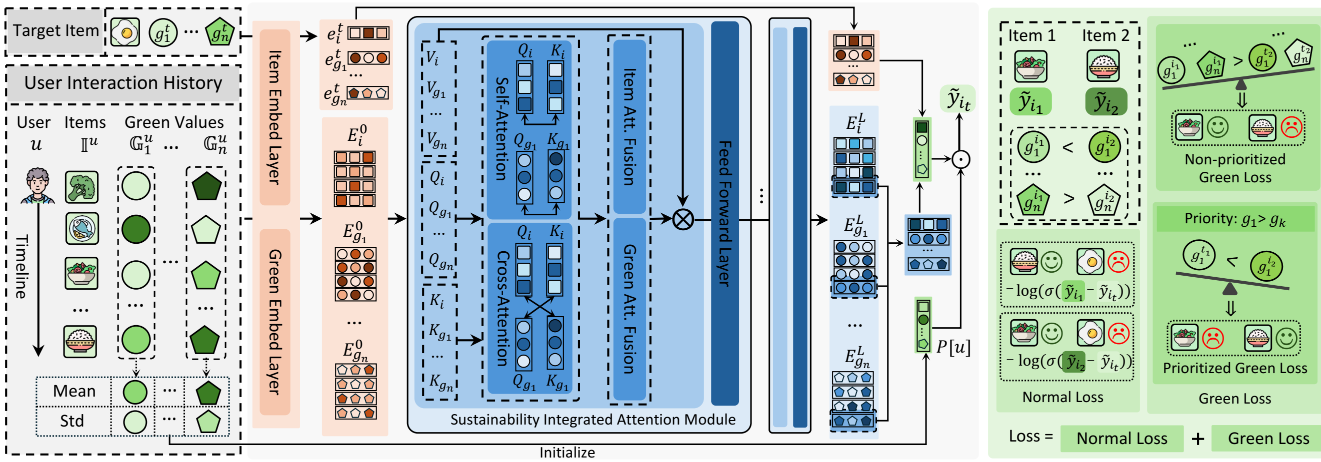

Figure 1: The model structure of *GRAPE*. *GRAPE* consists of three main modules: the Embedding Module, the Sustainability Integrated Attention Module, and the Prediction Module. The loss function is divided into two parts: the *Normal Loss*, which aims to enhance recommendation accuracy, and the *Green Loss*, which targets the improvement of recommendation sustainability. For the item's sustainability indicators, a deeper color signifies a greener value.

For the sustainability sequences $\mathbb{G}_1^u, ..., \mathbb{G}_n^u$, we introduce a joint embedding layer $\text{JE}(\cdot)$ to discretize and embed the sustainability indicator $g_j^{i_k}$:

$$\text{JE}(g_j^{i_k}) = \lfloor \frac{g_j^{i_k}}{\Delta} \rfloor \times \mathbf{L}_G, j \in \{1, ..., n\}, k \in \{1, ..., w\}, \quad (1)$$

where $\mathbf{L}_G \in \mathbb{R}^{\frac{max(g_j)}{\Delta} \times d}$, Δ is a hyperparameter that controls the level of granularity in the discretization process, $\lfloor \cdot \rfloor$ denotes the floor function (Chen, Yang, and Yu 2022). We then construct the embedding matrix $E_{g_1}^0, ..., E_{g_n}^0$ by concatenating the corresponding indicator embeddings.

Sustainability Integrated Attention Module Inspired by (Lin et al. 2024), we design a Sustainability Integrated Attention (SIA) Module to model the evolving preferences of users towards items and various sustainability indicators. Specifically, the SIA Module comprises four layers: a Projection Layer, Self- and Cross-Attention Layers, Attention Fusion Layers, and a Feed-Forward Layer. In *GRAPE*, we stack L SIA modules, and here we provide a detailed introduction to the l-th module.

Projection Layer: We first project each embedding matrix to capture query, key and value matrices with M attention heads. The m-th attention head for query matrix projection is formulated as:

$$Q_i^m = f_{Q_i^m}(E_i^{l-1}), Q_{g_j}^m = f_{Q_{g_j}^m}(E_{g_j}^{l-1}), j \in \{1, ..., n\}, \quad (2)$$

where $f_{Q_i^m}(\cdot), f_{Q_{g_j}^m}(\cdot)$ denote a linear projection layer, g_j denotes the j-th sustainability indicator. Similarly, we have key matrices $K_i^m, K_{g_j}^m$ and value matrices $V_i^m, V_{g_j}^m$.

Self- and Cross-Attention Layers: Given that users' preferences change over time and their behaviors demonstrate sequential dependencies, capturing these sequential behaviors is crucial for making accurate predictions (Jing et al. 2023; Chen, Yang, and Yu 2022). In *GRAPE*, we employ a self-attention mechanism to adeptly model users' evolving preferences for items and various sustainability indicators:

$$A_x^m = Q_x^m K_x^{m\top}, x \in \{i, g_1, ..., g_n\}. \quad (3)$$

Furthermore, we observe strong correlations among the items users interact with and their corresponding different sustainability indicators (Zhang et al. 2024a). For instance, users who frequently purchase nutritional food also tend to prefer items with high Healthy Meal Index scores. To capture such correlations within users' interaction sequences, we introduce a cross-attention layer:

$$A_{x,y}^m = Q_{x,y}^m K_{x,y}^{m\ \top}, x, y \in \{i, g_1, ..., g_n\}, x \neq y. \quad (4)$$

Attention Fusion layers: We then incorporate fusion layers to model user preferences towards items and various sustainability indicators. As for the item preference, we first aggregate all captured self- and cross-attention matrices:

$$\mathcal{A}_i = \{A_x, A_{y,z} | x, y, z \in \{i, g_1, ..., g_n\}, y \neq z\}. \quad (5)$$

Given $\mathcal{A}_i$, the Item Attention Fusion layer is formulated as

$$R_i^m = \text{Mask}(\text{Fuse}(\boldsymbol{r}_i, \mathcal{A}_i), V_i^m)), \quad (6)$$

where $\boldsymbol{r}_i \in \mathbb{R}^{(n+1)^2}$ is a learnable weight vector, $\text{Fuse}(\cdot)$ is a attention fusion function consists of weighted sum and gating operations (Liu et al. 2021), $\text{Mask}(\cdot)$ is a masked integration function which merges the captured attention matrix with value matrix V_i^m without information leakage (Kang and McAuley 2018).

Regarding user's sustainable preferences, we combine all attention matrices that include all sustainability-related information:

$$\mathcal{A}_g = \{A_x, A_{y,z} | x, y, z \in \{g_1, ..., g_n\}, y \neq z\}. \quad (7)$$

Then the embedding for each sustainability indicator preference is generated through the Green Attention Fusion layer:

$$R_{g_j}^m = \text{Mask}(\text{Fuse}(\boldsymbol{r}_g, \mathcal{A}_g), V_{g_j}^m)), j \in \{1, ..., n\}, \quad (8)$$

where $\boldsymbol{r}_g \in \mathbb{R}^{n^2}$ is a learnable weight vector.

Feed-Forward Layer: Finally, we leverage a Feed-Forward Network (FFN) to merge the output of all M attention heads (Vaswani 2017).

$$\begin{aligned} E_i^l &= \text{FFN}([R_i^1, ..., R_i^M]W_i), \\ E_{g_j}^l &= \text{FFN}([R_{g_j}^1, ..., R_{g_j}^M]W_{g_j}), j \in \{1, ..., n\}, \end{aligned} \tag{9}$$

where $W_i, W_{g_j} \in \mathbb{R}^{d \times d}$ are learnable matrices.

We leverage L SIA modules. For the l-th SIA module, the matrices $E_i^l, E_{g_j}^l (j \in \{1, ..., n\})$ are then defined as:

$$E_i^l, E_{g_1}^l, ..., E_{g_n}^l = \text{SIA}_l(E_i^{l-1}, E_{g_1}^{l-1}, ..., E_{g_n}^{l-1}), \tag{10}$$

where $\text{SIA}_l, l \in \{1, ..., L\}$ denotes the l-th SIA module. We then extract the last embedding vector from each matrix as the final preference representations, denoted as $\boldsymbol{e}_i^L, \boldsymbol{e}_{g_j}^L$, respectively (Devlin 2018).

Prediction Module Different users have varying preferences and levels of acceptance for sustainable food. Some might prefer junk food but also open to trying healthy options occasionally while some others may strictly favor junk food and show significant resistance to changing their dietary habits. In this module, we employ an attention layer to explicitly model users' preferences for items and their inclination to select items with high sustainability values. Specifically, we first multiply each output of SIA modules $\boldsymbol{e}_i^L, \boldsymbol{e}_{g_1}^L, ..., \boldsymbol{e}_{g_n}^L$ with the embeddings $\boldsymbol{e}_i^0, \boldsymbol{e}_{g_1}^0, ..., \boldsymbol{e}_{g_n}^0$ obtained by Embedding Module, respectively. Then, we introduce an attention matrix $P \in \mathbb{R}^{v \times (n+1)}$:

$$\widetilde{y} = [\boldsymbol{e}_i^L \boldsymbol{e}_i^{0\top}, \boldsymbol{e}_{g_1}^L \boldsymbol{e}_{g_1}^{0\top}, ..., \boldsymbol{e}_{g_n}^L \boldsymbol{e}_{g_n}^{0\top}]^\top \times P[u], \tag{11}$$

where $P[u] \in \mathbb{R}^{1 \times (n+1)}$ is the u-th vector in P, representing u's preference for the item itself and various sustainability dimensions.

Training

Initialization We observe that the average values of sustainability indicators for items users have previously interacted with can somewhat reflect their preferences for sustainable food (Zhang et al. 2024b). A higher average usually indicates that the user prefer more sustainable food. Similarly, the standard deviation values of these indicators shed light on their acceptance. When users' daily eating habits are both relatively unhealthy and consistent, a larger variance may suggest a greater likelihood of users adopting more sustainable options. To make full use of the users' interaction history and improve training, we propose to initialize the attention matrix P using:

$$P[u] = [1, \mu(\mathbb{G}_1^u)\lambda(\mathbb{G}_1^u), ..., \mu(\mathbb{G}_n^u) \cdot \lambda(\mathbb{G}_n^u)], \tag{12}$$

where $\mu(\cdot)$ and $\lambda(\cdot)$ represent the functions for calculating the mean and variance of a sustainability indicator sequence, respectively. $P[u]$ is regularize to meet $\sum P[u] = 1$.

Sampling and Loss Functions

Given a user u, we use $\mathcal{I}_u^+, \mathcal{I}_u^-$ to denote the sets of items that user u has interacted with and not interacted with, respectively. The loss $\mathcal{L}$ contains two components: *Normal Loss* $\mathcal{L}_n$ which helps *GRAPE* generate accurate recommendations based on the user's interaction history, and *Green Loss* $\mathcal{L}_g$, designed to enhance the visibility of more sustainable food options:

$$\mathcal{L} = \alpha \mathcal{L}_n + (1 - \alpha)\mathcal{L}_g, \tag{13}$$

where $\alpha \in [0, 1]$ is a hyperparameter.

As for the *Normal Loss* $\mathcal{L}_n$, we sample $i^+ \in \mathcal{I}_u^+$ and $i^- \in \mathcal{I}_u^-$ and leverage the pair wise BPR Loss (Rendle et al. 2012):

$$\mathcal{L}_n = -\sum_{<u, i^+, i^->} log(\sigma(\widetilde{y}_{i^+} - \widetilde{y}_{i^-})), \tag{14}$$

where $\widetilde{y}_i$ denotes the predicted probability that user u will interact with item i as determined by *GRAPE*, σ refers to the sigmoid function.

The *Green Loss* $\mathcal{L}_g$ is designed to assign higher scores to more sustainable items. However, it's crucial to align these scores with user preferences, rather than merely increasing the visibility of greener products. A recommendation is considered unsuccessful if a user dislikes the recommended item, regardless of its sustainability. Therefore, we differentiate between items the user has interacted with, $\mathcal{I}_u^+$, and those they haven't, $\mathcal{I}_u^-$, to ensure that our recommendations accurately reflect both sustainability and user preferences. Specifically, we sample $i_1, i_2 \in \mathcal{I}_u^+$ or $i_1, i_2 \in \mathcal{I}_u^-$, where $[g_1^{i_1}, ..., g_n^{i_1}]$ represents the sustainability indicator values for i_1, and $[g_1^{i_2}, ..., g_n^{i_2}]$ represents those for i_2. We assume that a higher value of g_j^i indicates that item i is more sustainable according to the j-th sustainability indicator. For sustainability indicators where a smaller value signifies greater sustainability, we apply inverse conditions when calculating the loss.

We introduce two types of *Green Loss* functions: the *Non-prioritized Green Loss* $\mathcal{L}_{ng}$ treats all sustainability indicators equally, while the *Prioritized Green Loss* $\mathcal{L}_{pg}$ assigns different priorities to each indicator. The *Non-prioritized Green Loss* $\mathcal{L}_{ng}$ is then defined as:

$$\mathcal{L}_{ng} = -\sum_{j=1}^{n} log(\sigma((g_j^{i_1} - g_j^{i_2})(\widetilde{y}_{i_1} - \widetilde{y}_{i_2}))). \tag{15}$$

The *Prioritized Green Loss* $\mathcal{L}_{pg}$ is designed to recognize the varying significance of different sustainability indicators. For example, in food recommendations, factors like health and nutrition might be considered more crucial than the environmental impact of food production. Therefore, $\mathcal{L}_{pg}$ is formulated to incorporate these manually defined priorities. Inspired by (Zhang et al. 2023), we propose considering an indicator only if all the indicator values with higher priorities exceed predefined thresholds. Specifically, assuming a priority order of $g_1 > g_2 > \cdots > g_n$, we first compare the values of g_1 for items i_1 and i_2. An indicator g_j is then considered valid if, for all g_k with $k < j$, the condition $min(g_k^{i_1}, g_k^{i_2}) \geq \beta_k$ is satisfied, where β_k represents the threshold for the indicator g_k. $\mathcal{L}_{pg}$ is defined as:

$$\mathcal{L}_{pg} = -\sum_{j=1}^{n} D(j) \cdot log(\sigma((g_j^{i_1} - g_j^{i_2})(\widetilde{y}_{i_1} - \widetilde{y}_{i_2}))). \tag{16}$$

where

$$D(j) = \begin{cases} 1 & \text{if } j = 1 \text{ and } min(g_1^{i_1}, g_1^{i_2}) < \beta_1 \\ 1 & \text{if } j > 1 \text{ and } \begin{array}{l} min(g_j^{i_1}, g_j^{i_2}) < \beta_j; \\ min(g_k^{i_1}, g_k^{i_2}) \geq \beta_k, \\ \forall k \in \{1, ..., j-1\} \end{array} \\ 0 & \text{else} \end{cases} \quad (17)$$

Experiment Settings

Dataset

To evaluate the effectiveness of *GRAPE*, we use a food recommendation dataset (Zhang et al. 2024a) which contains 6290 users, 74324 recipes, and 316116 interactions. Each user or recipe has at least 10 interactions. Each recipe is labeled with three widely used sustainability indicators: 1) Environmental Impact Score (EIS) (Clark et al. 2022; Gephart et al. 2021), where a lower score indicates greater environmental friendliness; 2) Nutritional Impact Score (NIS) (Clark et al. 2022), with a higher NIS score denoting more nutritious food; and 3) Healthy Meal Index (HMI) (Kasper et al. 2016), where a higher HMI score suggests a healthier meal. For dataset splitting, we employ a leave-one-out method.

Baselines

We select the following baselines:

- **BPR** (Rendle et al. 2012), a traditional model that ranks items by maximizing the predicted score difference between a user's preferred and non-preferred items. (Wang, De Vries, and Reinders 2006), a collaborative filtering method that recommends items by selecting items similar to user previously interacted items.
- **SHT** (Xia, Huang, and Zhang 2022), a non-sequential baseline that enhances conventional matrix factorization with a hypergraph transformer network and generative self-supervised data augmentation .
- **STOSA** (Fan et al. 2022), which embeds each item as a stochastic Gaussian distribution, and forecasts the next item with a self-attention mechanism.
- **ICLRec** (Chen et al. 2022), which learns users' behaviors from unlabeled user historical interactions and is optimized through contrastive self-supervised learning.
- **NOVA** (Liu et al. 2021), which introduces a non-invasive attention mechanism to inject side-information into BERT structure for better attention distribution.
- **CAFE** (Li et al. 2022), which explicitly models user preference by fusing fine-grained item representations and coarse-grained side-information representations.
- **FDSA-CL** (Hao et al. 2023), which introduces independent self-attention layers for item and feature representations and leverage contrastive method for training.
- **MSSR** (Lin et al. 2024), which leverage self-attention machinism to capture item-feature and feature-feature correlation for both item and side-information prediction.
- **FHFRS** (Rostami, Aliannejadi, and Oussalah 2023), which is a post-process green food recommendation model, and we adopt **MSSR** as its framework.

Implementation Details

In our experiment, we select learning rates from [0.0001, 0.01], embedding size d from{32, 64}, and batch size from {32, 64, 128, 256}. L2 regularization is applied when computing the loss function. For hyperparameters, we select α from $[0.5, 1]$ for the loss function. In prioritized green loss, we select hyperparameter β_{EIS} from $[70, 120]$, β_{NIS} from $[30, 50]$, and β_{HMI} from $[30, 50]$. [1]

Evaluation Metrics

To evaluate the performances of Top-N sequential recommendation, we leverage the hit ratio (HR) and normalized discounted cumulative gain (NDCG) for $N = 5, 10, 20$. We also calculate the average values of each sustainability indicator in the recommended list.

Results and Analysis

Overall Performance

Table 1 demonstrates the top-N recommendation performances of *GRAPE* and all baselines. From the results, we make the following observations. First, in terms of recommendation accuracy, *GRAPE* consistently attains the top position across all NDCG and most HR metrics. For instance, *GRAPE* leads the best baseline MSSR by 4.18% in NDCG@5 and 3.94% in NDCG@20. These findings demonstrate that *GRAPE* can effectively model user preferences and suggest that integrating sustainability indicators may further improve recommendation accuracy.

Second, the food recommended by *GRAPE* consistently achieves higher sustainability scores than those recommended by baseline models, especially against state-of-the-art models with notable top-N recommendation accuracy. Specifically, in terms of the Environmental Impact Score (EIS), *GRAPE* outperforms MSSR by 8.91% at N=10 and by 10.84% at N=20. This highlights *GRAPE*'s ability to adaptively recommend sustainable items that align with user preferences. While traditional methods like KNN also recommend items with relatively high sustainability scores, they tend to achieve lower recommendation accuracy.

In short, our *GRAPE* outperforms all state-of-the-art baselines in both recommendation accuracy and sustainability indicators. This demonstrates that recommending greener foods does not compromise the accuracy of recommendations and can, in fact, enhance them simultaneously.

Effectiveness of Green Loss

To assess the effectiveness of our proposed Green Loss, we perform an ablation study using the Non-prioritized Green Loss on the hyperparameter α in Equation 13. A higher α value indicates a lower proportion of green loss. The results are shown in Figure 2a.

Consistent with our assumption, as α increases, the recommendation list consistently displays lower values across all sustainability indicators. This suggests that our proposed green loss can effectively encourage the model to make greener recommendations.

[1] Our code is available: https://github.com/JingXiaoyi/GRAPE.

Top-N	Metrics	BPR	KNN	SHT	STOSA	ICLRec	NOVA	CAFE	FDSA-CL	FHFRS	MSSR	*GRAPE* ($\mathcal{L}_{np}$)
N=5	HR	0.0285	0.0461	0.0428	0.0663	0.0618	0.0703	0.0634	0.0681	0.0670	0.0729	**0.0738**
	NDCG	0.0116	0.0195	0.0191	0.0298	0.0316	0.0341	0.0314	0.0327	0.0322	0.0359	**0.0374**
	EIS↑	87.94	**81.54**	82.57	105.49	102.64	112.68	103.19	107.83	106.67	100.47	95.95
	NIS↑	33.26	32.21	27.84	33.69	**34.61**	27.51	30.16	29.84	33.89	33.53	34.26
	HMI	41.55	42.72	44.16	43.66	41.84	42.69	43.08	41.76	43.76	42.16	**44.17**
N-10	HR	0.0307	0.0629	0.0614	0.0984	0.1012	0.1131	0.0998	0.1107	0.1140	**0.1184**	0.1152
	NDCG	0.0123	0.0247	0.0252	0.0401	0.0454	0.049	0.0431	0.0473	0.0489	0.0512	**0.0519**
	EIS↑	89.00	**84.42**	93.25	100.97	95.57	101.45	97.26	102.07	96.99	100.18	91.25
	NIS↑	32.16	32.52	28.13	33.23	27.10	29.71	31.91	32.46	33.79	33.97	**34.77**
	HMI	43.72	42.84	42.23	**44.52**	42.32	41.26	41.59	43.81	42.70	42.86	44.19
N=20	HR	0.0433	0.0952	0.0997	0.1584	0.1416	0.1615	0.1572	0.1593	0.1539	0.1657	**0.1665**
	NDCG	0.0155	0.0324	0.0349	0.0563	0.0557	0.0609	0.0571	0.0594	0.0581	0.0634	**0.0659**
	EIS↑	87.35	**80.42**	85.23	96.10	98.27	97.28	96.28	94.17	94.51	97.32	86.77
	NIS↑	32.96	32.11	24.10	33.28	29.50	30.58	31.85	33.28	32.67	34.15	**34.56**
	HMI	43.17	42.37	44.21	42.50	41.82	42.03	41.72	43.81	43.24	42.65	**44.30**

Table 1: Performances of different methods for Top-N recommendation. The best results are bold, and the second-best are underlined. A lower EIS indicates greater environmental friendliness, whereas higher NIS and HMI values denote more nutritious and healthier food, respectively.

We also observe that as α increases, the accuracy of recommendations first rises rapidly and then begins to fluctuate, with peak performance occurring at $\alpha = 0.9$. These results highlight the need to balance the model's focus on user preferences and item sustainability. Although making greener recommendations contributes to higher recommendation accuracy, prioritizing users' personal preferences remains paramount to the recommendation process.

Effectiveness of User Attention Matrix

In this work, we assume that users have varying preferences and acceptance levels for sustainable items. To address this, we design a user attention matrix P to explicitly capture these preferences. We initialize P based on insights derived from users' historical interactions. To validate our assumption and assess the effectiveness of our proposed initialization methods, we implement several variants of P using different initialization approaches:

- *Fixed* P_{one}: We replace P with frozen all-ones matrix $P_{one} = \mathbf{1}^{v\times(n+1)}$, indicating that there are no distinctions in personal preferences among items or sustainability indicators.
- P_N + *RandInit*: We apply a consistent vector $\boldsymbol{p} \in \mathbb{R}^{(n+1)}$ to populate each row of the matrix $P_N \in \mathbb{R}^{v\times(n+1)}$. $\boldsymbol{p}$ is randomly initialized. This variant indicates that all users share a uniform preference towards sustainability.
- P_N: Similar to P_N + *RandInit*, we apply a consistent vector $\boldsymbol{p} \in \mathbb{R}^{(n+1)}$ to represent all users's preferences. We then initialize $\boldsymbol{p}$ using the average and standard deviation of the sustainability values from items interacted with by all users.
- P + *RandInit*: We use randomly initialization for the user attention matrix P.
- $P_{\textit{GRAPE}}$: This is the default setting for *GRAPE*.

The results are illustrated in Figure 2b. First, we observe that Fixed P_{one} performs significantly worse than all other variations, supporting our assumption that users have different preferences among items or sustainability indicators. Moreover, when comparing the non-personalized variations, P_N + *RandInit* and P_N, with the personalized ones, P + *RandInit* and $P_{\textit{GRAPE}}$, we find personalized group performs better. This suggests that users' preferences for different attributes vary from person to person. Furthermore, $P_{\textit{GRAPE}}$ outperforms all other variations, demonstrating the effectiveness of initialization using the average and standard deviation of sustainability values from users' interacted items.

Model	HR	NDCG	EIS↓	NIS↑	HMI↑
GRAPE ($\mathcal{L}_{pg}$)	0.1122	0.0513	●**86.30**	◐32.15	○42.08
	0.1129	0.0512	●87.74	○32.06	◐41.93
	0.1132	0.0519	◐93.26	●35.13	○41.52
	0.1142	0.0513	○94.22	●**35.72**	◐44.16
	0.1133	0.0521	◐88.54	○33.04	●44.95
	0.1138	**0.0525**	○93.32	◐35.22	●**45.26**
GRAPE ($\mathcal{L}_{ng}$)	**0.1152**	0.0519	91.25	34.77	44.19

Table 2: Top-10 performance of *GRAPE* with *Non-prioritized Green Loss* ($\mathcal{L}_{ng}$), and *Prioritized Green Loss* ($\mathcal{L}_{pg}$) applying different priority orders. ● denotes the highest priority, ◐ denotes the second priority, and ○ denotes the lowest priority. The best results for each evaluation metric are bold.

Ablation Study for Prioritized Green Loss

Our prioritized green loss $\mathcal{L}_{pg}$ allows *GRAPE* to make greener recommendations by prioritizing sustainability indicators in a specific order. To evaluate the effectiveness of $\mathcal{L}_{pg}$, we experiment with various priority orders and set different thresholds β for the indicators within $\mathcal{L}_{pg}$.

Table 2 displays the performances when using different priority orders for three sustainability indicators. In line with our expectations, the sustainability indicator assigned the highest priority consistently achieves the best performance

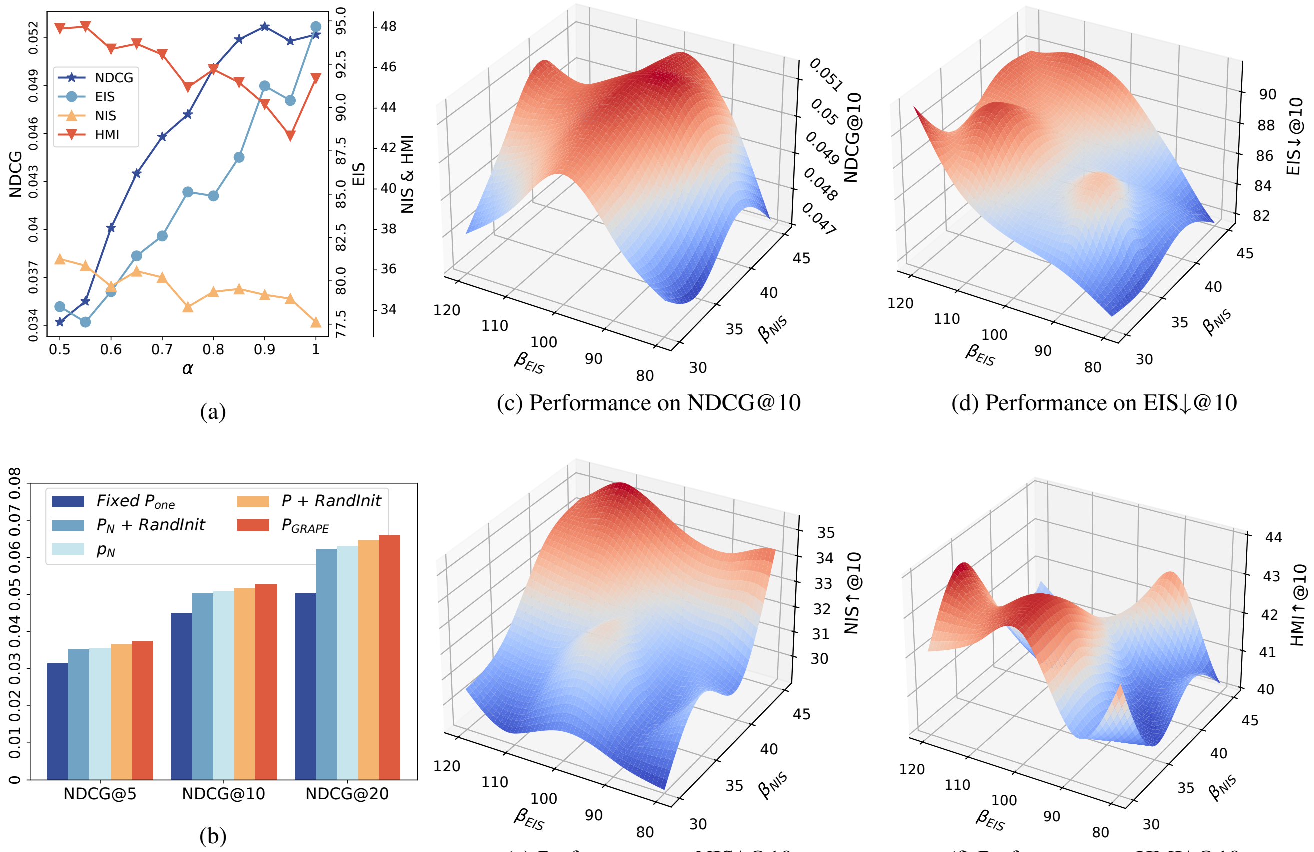

(a)

(c) Performance on NDCG@10

(d) Performance on EIS↓@10

(b)

(e) Performance on NIS↑@10

(f) Performance on HMI↑@10

Figure 2: Ablation studies: (a) Top-10 performances of *GRAPE* with different hyperparameter α for the *Non-prioritized Green Loss*. (b) Performances of *GRAPE* with different user attention matrix. (c)-(f) Performances of recommendation accuracy and sustainability under varying β_{EIS} and β_{NIS} in the *Prioritized Green Loss*, with the priority sequence set as EIS>NIS>HMI.

within its category. Additionally, we observe that setting either NIS or HMI as the highest priority results in relatively high performance for the other indicator as well. This may be due to a high correlation between the food nutritional level and the healthy meal index.

Then we select the priority EIS>NIS>HMI and conduct ablation study on the sustainability thresholds β_{EIS} and β_{NIS}. It is important to note that a smaller β_{EIS} or a larger β_{NIS} imposes stricter constraints on the corresponding sustainability indicator, which is expected to yield more sustainable recommendations within that category. The results are shown in Figure 2c-2f, with each graph representing a different evaluation metric.

As shown in Figure 2d and 2e, we observed that when the constraints on a specific indicator are tightened, the performances of the corresponding indicator improves generally. Such findings demonstrate that our loss function effectively adjusts the model's focus across different indicators.

From Figure 2c, we observe that the recommendation accuracy first increases and then decreases with the rise in either β_{EIS} or β_{NIS}. This may be attributed to the model's tendency to focus primarily on a particular emphasized sustainability indicator, which can sometimes come at the expense of overall recommendation performance. It underscores the necessity to strike a balance between user preferences and item sustainability to optimize recommendation effectiveness.

Conclusion

In conclusion, recommendation systems have significant potential to encourage more sustainable choices among users. However, most existing methods focus solely on recommendation accuracy, often overlooking the potential impact of recommending more sustainable items. To address this oversight, we introduce the Green Food Recommendation task and propose a novel method called *Green Recommender Aligned with Personalized Eating* (*GRAPE*). *GRAPE* not only models users' evolving preferences for items but also their willingness to choose sustainable foods. Extensive experiments demonstrate the superiority of *GRAPE*, showing that it successfully balances recommendation accuracy with enhanced sustainability attributes of the recommended foods. We believe that our research lays a strong foundation to encourage future studies in this field.

Acknowledgments

This research is supported by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore through Alibaba-NTU Global e-Sustainability CorpLab (ANGEL). It is also supported by the Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY) and College of Computing and Data Science (CCDS) at NTU Singapore.

References

Chen, L.; Yang, N.; and Yu, P. S. 2022. Time lag aware sequential recommendation. In *CIKM*, 212–221.

Chen, Y.; Liu, Z.; Li, J.; McAuley, J.; and Xiong, C. 2022. Intent contrastive learning for sequential recommendation. In *Proceedings of the ACM Web Conference 2022*, 2172–2182.

Clark, M.; Springmann, M.; Rayner, M.; Scarborough, P.; Hill, J.; Tilman, D.; Macdiarmid, J. I.; Fanzo, J.; Bandy, L.; and Harrington, R. A. 2022. Estimating the environmental impacts of 57,000 food products. *Proceedings of the National Academy of Sciences*, 119(33): e2120584119.

Curto, D.; Franzitta, V.; and Guercio, A. 2021. A review of the water desalination technologies. *Applied Sciences*, 11(2): 670.

Darvishi, P.; Mousavi, S. A.; Mahmoudi, A.; and Nayeri, D. 2023. A comprehensive review on the removal of antibiotics from water and wastewater using carbon nanotubes: synthesis, performance, and future challenges. *Environmental Science: Water Research & Technology*, 9(1): 11–37.

Devlin, J. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Fan, Z.; Liu, Z.; Wang, Y.; Wang, A.; Nazari, Z.; Zheng, L.; Peng, H.; and Yu, P. S. 2022. Sequential recommendation via stochastic self-attention. In *Proceedings of the ACM Web Conference 2022*, 2036–2047.

Fang, W.; and Durable, C. 2024. carbon dioxide efficiently to reuse and recycle it. *Nature*, 626: 86–91.

Forouzandeh, S.; Berahmand, K.; Rostami, M.; Aminzadeh, A.; and Oussalah, M. 2024. UIFRS-HAN: User interests-aware food recommender system based on the heterogeneous attention network. *Engineering Applications of Artificial Intelligence*, 135: 108766.

Gephart, J. A.; Henriksson, P. J.; Parker, R. W.; Shepon, A.; Gorospe, K. D.; Bergman, K.; Eshel, G.; Golden, C. D.; Halpern, B. S.; Hornborg, S.; et al. 2021. Environmental performance of blue foods. *Nature*, 597(7876): 360–365.

Hao, Y.; Zhang, T.; Zhao, P.; Liu, Y.; Sheng, V. S.; Xu, J.; Liu, G.; and Zhou, X. 2023. Feature-level deeper self-attention network with contrastive learning for sequential recommendation. *IEEE transactions on knowledge and data engineering*, 35(10): 10112–10124.

He, X.; Liu, Q.; and Jung, S. 2024. The impact of recommendation system on user satisfaction: A moderated mediation approach. *Journal of Theoretical and Applied Electronic Commerce Research*, 19(1): 448–466.

Jain, P.; Castellanos-Acuna, D.; Coogan, S. C.; Abatzoglou, J. T.; and Flannigan, M. D. 2022. Observed increases in extreme fire weather driven by atmospheric humidity and temperature. *Nature Climate Change*, 12(1): 63–70.

Jing, J.; Zhang, Y.; Zhou, X.; and Shen, Z. 2023. Capturing Popularity Trends: A Simplistic Non-Personalized Approach for Enhanced Item Recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 1014–1024.

Jones, N. 2017. How machine learning could help to improve climate forecasts. *Nature*, 548(7668).

Kang, W.-C.; and McAuley, J. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*, 197–206. IEEE.

Kasper, N.; Mandell, C.; Ball, S.; Miller, A. L.; Lumeng, J.; and Peterson, K. E. 2016. The Healthy Meal Index: A tool for measuring the healthfulness of meals served to children. *Appetite*, 103: 54–63.

Kirby, C. K.; and Zwickle, A. 2021. Sustainability behaviors, attitudes, and knowledge: Comparing university students and the general public. *Journal of Environmental Studies and Sciences*, 11(4): 639–647.

Li, J.; Zhao, T.; Li, J.; Chan, J.; Faloutsos, C.; Karypis, G.; Pantel, S.-M.; and McAuley, J. 2022. Coarse-to-fine sparse sequential recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 2082–2086.

Lin, X.; Luo, J.; Pan, J.; Pan, W.; Ming, Z.; Liu, X.; Huang, S.; and Jiang, J. 2024. Multi-Sequence Attentive User Representation Learning for Side-information Integrated Sequential Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 414–423.

Liu, C.; Li, X.; Cai, G.; Dong, Z.; Zhu, H.; and Shang, L. 2021. Noninvasive self-attention for side information fusion in sequential recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 4249–4256.

Marcus, J.; and Roy, J. 2019. In search of sustainable behaviour: The role of core values and personality traits. *Journal of Business Ethics*, 158: 63–79.

Ngcamu, B. S. 2023. Climate change effects on vulnerable populations in the Global South: a systematic review. *Natural Hazards*, 118(2): 977–991.

Orlinska-Wozniak, P.; Szalinska, E.; Jakusik, E.; Bojanowski, D.; and Wilk, P. 2021. Biomass production potential in a river under climate change scenarios. *Environmental Science & Technology*, 55(16): 11113–11124.

Peleg Mizrachi, M.; and Tal, A. 2022. Regulation for promoting sustainable, fair and circular fashion. *Sustainability*, 14(1): 502.

Poore, J.; and Nemecek, T. 2018. Reducing food's environmental impacts through producers and consumers. *Science*, 360(6392): 987–992.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*.

Rostami, M.; Aliannejadi, M.; and Oussalah, M. 2023. Towards health-aware fairness in food recipe recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*, 1184–1189.

Ruggerio, C. A. 2021. Sustainability and sustainable development: A review of principles and definitions. *Science of the Total Environment*, 786: 147481.

Samir, A.; Ashour, F. H.; Hakim, A. A.; and Bassyouni, M. 2022. Recent advances in biodegradable polymers for sustainable applications. *Npj Materials Degradation*, 6(1): 68.

Schwartz, R.; Dodge, J.; Smith, N. A.; and Etzioni, O. 2020. Green ai. *Communications of the ACM*, 63(12): 54–63.

To, A.; Pham, H. Q. V.; Nguyen, Q. H.; Davis, J. G.; O'Sullivan, B.; Pan, S. L.; and Nguyen, H. D. 2024. Carbon Stock Estimation at Scale from Aerial and Satellite Imagery. In *2024 IEEE Conference on Artificial Intelligence (CAI)*.

Vaswani, A. 2017. Attention is All You Need. *arXiv preprint arXiv:1706.03762*.

Verdecchia, R.; Sallou, J.; and Cruz, L. 2023. A systematic review of Green AI. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(4): e1507.

Štofejová, L.; Král', S.; Fedorko, R.; Bačík, R.; and Tomášová, M. 2023. Sustainability and Consumer Behavior in Electronic Commerce. *Sustainability*, 15(22).

Wang, J.; De Vries, A. P.; and Reinders, M. J. 2006. Unifying user-based and item-based collaborative filtering approaches by similarity fusion. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, 501–508.

Xia, L.; Huang, C.; and Zhang, C. 2022. Self-supervised hypergraph transformer for recommender systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2100–2109.

Yin, G.; Guan, P.; Wang, Y.-H.; Zhang, P.; Qu, B.; Xu, S.; Zhang, G.; He, C.; Shi, Q.; and Wang, J. 2024. Temporal Variations in Fire Impacts on Characteristics and Composition of Soil-Derived Dissolved Organic Matter at Qipan Mountain, China. *Environmental Science & Technology*.

Zhang, L.; Zhang, Y.; Zhou, X.; and Shen, Z. 2024a. GreenRec: A Large-Scale Dataset for Green Food Recommendation. In *Companion Proceedings of the ACM on Web Conference 2024*, 625–628.

Zhang, S.; Jia, F.; Wang, C.; and Wu, Q. 2023. Targeted hyperparameter optimization with lexicographic preferences over multiple objectives. In *The Eleventh international conference on learning representations*.

Zhang, Y.; Zhou, X.; Meng, Q.; Zhu, F.; Xu, Y.; Shen, Z.; and Cui, L. 2024b. Multi-modal Food Recommendation using Clustering and Self-supervised Learning. *arXiv preprint arXiv:2406.18962*.