International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 11 Issue: 05 | May 2024 www.irjet.net p-ISSN: 2395-0072

Recipe Renaissance: Leveraging Deep Reinforcement Learning for Food Recommendations Systems in Culinary Innovation.

Neha Tyagi[1], Vijay Maheshwari[2]

[1]M.tech [Department of Computer Science &Engineering] Shobhit University, Meerut .Uttar Pradesh.
[2]Associate Professor [Department of Computer Science &Engineering] Shobhit University, Meerut .Uttar Pradesh.

---***---

Abstract - *Personalized meal planning systems are important in addressing individual nutritional needs, health concerns, and dietary choices. However, challenges exist when effectively accommodating diverse user requirements, resulting in low user preference adherence and satisfaction. This research suggests a novel approach that combines the power of collaborative filtering (CF) and deep reinforcement learning (DRL) by applying the Deep Deterministic Policy Gradient (DDPG) algorithm to effectively handle these obstacles completely. Nutrition, health, preferences, culture, flavor, choices, and dietary restrictions factors like these usually influence eating habits. Traditional systems find difficulties in matching meal suggestions with the user's lifestyle, preferences, and individual dietary requirements or restrictions, owning to less-than-ideal user interactions. Our framework leverages CF insights into user preferences, and dietary restrictions, and then applies DRL's adaptive decision-making based on users' interactions as feedback. Using the DDPG algorithm to integrate CF with DRL, our approach effectively caters to most of the user requirements and preferences. DRL principles like Markov decision processes (MDPs) are essential for interactive meal recommendations. A reward-shaping feature is included offering multi-criteria decision-making in producing personalized meal plans for the user. Through this research study, we've compared the effectiveness of our proposed integrated methodology against conventional methods. This study sheds light on the possibility of CF-DRL integration in providing individualized meal ideas, promising to improve user experiences and health goals in the best possible way.*

Key Words: personalized meal planning, collaborative filtering, deep reinforcement learning, DDPG algorithm, user preferences, dietary choices, nutritional needs, health considerations.

1. INTRODUCTION

In the present digital world, we are living in, the intermixing of culinary exploration and technological advancements has resulted in numerous food recipe platforms, therefore people engage with them differently. Central to this culinary evolution is the creation and utilization of integrated food recommendation systems, seamlessly mixing the domains of flavor and health to provide customized culinary journeys that cater to personal tastes and dietary requirements. The analysis of all these interrelated elements shows that creating the most effective customized meal plans is important as well as challenging. Despite their emphasis on meeting nutritional and health needs, traditional meal planning systems have never been able to accommodate people's diverse tastes, lifestyles, or covert eating patterns [1][2][3]. As there is a disconnect in between what we expect and what the system performs often results in partial or no user requirements adherence, thereby affecting the efficacy of the solutions systems are producing [1][4]. Therefore, personalized meal planning systems are undergoing evolution to better accommodate these diverse needs day and night. The objective for our research project is provided by this introduction, which tells how we can predict the interdependent variables between user satisfaction, following suggested meal plans, and the incorporation of user preferences and nutritional data into these creative systems. This study reveals the details of customized food recommendations and the resulting benefits for encouraging healthier eating habits and improving general well-being by diving into the complicated phenomena of nutritional value and taste.

Customized nutrition that considers an individual's needs, interests, and health circumstances is what personalized meal planning is all about—it's not simply a cuisine craze. Still, there are obstacles to overcome to achieve this degree of customization. When it comes to delivering consumers with recommendations that are too generic or generic, traditional meal-planning systems frequently fall short. The result is that it is difficult for consumers to adhere to these customized strategies, which lowers engagement rates. Though customized nutrition makes a lot of promises, it is frequently difficult to get the intended results.

To address these obstacles, this research presents an innovative solution that combines collaborative filtering (CF) with deep reinforcement learning (DRL) through the utilization of the Deep Deterministic Policy Gradient (DDPG) algorithm. The combination results in a significant advancement in personalized meal preparation, providing a comprehensive structure that seamlessly integrates data-driven analysis with flexible decision-making processes.

Our proposed methodology gains useful insights into users' cultural practices, preferences, and dietary restrictions using the basic approach which is collaborative Filtering. These insights serve as the foundation upon which our Deep

Learning Filtering model functions, thereby dynamically adjusting meal suggestions based on user feedback and real-time interactions likes, dislikes, ratings, etc. Through the complementary strengths of CF and DRL, we aim to bridge the gap between generic meal plans and truly customized nutrition solutions.

In the following sections of the study, we will give a literature review of all food recommendation systems that have been researched so far, after that will explain the intricacies of our integrated CF-DRL approach, including its key components, methodologies, and its impact on customized meal planning. By performing practical assessment and comparative research, we try to showcase the effectiveness of our proposed algorithm in enhancing user satisfaction, and adherence levels, and meeting general health and wellness goals.

2. Literature survey

Recommender systems serve a vital role in delivering customized recommendations across diverse fields to tackle the challenge of information overload. These systems are specifically crafted to offer tailored suggestions aligned with individual user preferences, thereby enriching user experience and facilitating decision-making processes. The rise of customized meal planning platforms has garnered considerable interest from a variety of disciplines, such as nutrition science, computer science, and human-computer interaction [1, 5,6]. We explore a range of methods employed in the quest for effective meal-planning solutions, ranging from rule-based systems to advanced machine learning (ML) approaches.

While prior studies have predominantly concentrated on creating computational instruments for dietary guidance, also they have not fully incorporated user preferences and nutritional data [5]. Raciel Yera Toledo et.al proposed framework aims to address the research gap by simultaneously considering user preferences and nutritional information in daily meal plan recommendations. The research endeavors to address this gap comprehensively by proposing a framework that generates personalized daily meal plans accommodating both user preferences and nutritional requirements.

2.1 Conventional Methods for Food Recommendation Systems

Rule-based systems employ predefined rules and constraints based on nutritional guidelines to produce meal plans, lacking flexibility for individual preferences or lifestyle changes [12,13]

Expert systems utilize nutritionist expertise to create personalized meal plans but require frequent updates to stay current with nutritional research and guidelines. [12,13]

Traditionally, food recommendation systems have focused on providing personalized food or recipe suggestions to individual users, often overlooking group interactions [6]. These systems cannot frequently offer recommendations tailored for group activities such as dining out, neglecting an important aspect of user preferences. To address this gap researchers like Mehrdad Rostami et.al have explored community detection techniques to identify user groups with similar tastes, enhancing recommendation accuracy. Time-aware algorithms are essential, taking into account the temporal dimension of user engagements to maintain the relevance of recommendations. Group recommendation systems are more complex than individual ones, as they must cater to the preferences of multiple users within a group. The proposed model in this study integrates metrics for group satisfaction, contributing to an improved recommendation process. Their study presents an innovative group food recommendation system that utilizes deep social community detection and user popularity, aiming to overcome limitations identified in current models. By incorporating features from food recommenders, community detection, time-aware algorithms, and group recommenders, their proposed system demonstrates superior performance compared to existing literature.

Many research initiatives have delved into the development of prediction models for type-2 diabetes harnessing the power of artificial intelligence (AI), machine learning, and deep learning methodologies. The primary objective has been to bolster early detection capabilities and provide tailored health recommendations for individuals. [6] It creatively employs a user-cloud environment for real-time applications, demonstrating the DP-UCE framework's ensemble of machine learning models and cloud-based user applications for type-2 diabetes prediction and suggesting diet plans. Moreover, recent research has explored innovative methodologies such as hybrid stacked ensemble merging genetic algorithms, healthcare monitoring systems reliant on convolutional neural networks (CNNs), and fused machine learning models. These efforts are geared towards enhancing predictive accuracy and enabling real-time monitoring in the management of diabetes. The integration of Internet of Things (IoT) technology, deep learning algorithms, and ensemble techniques signals a notable shift towards more comprehensive and efficacious strategies for diabetes prediction and management evident in contemporary literature.

Machine learning is rapidly becoming one of the most dynamic technologies applied in healthcare. A range of machine learning techniques is being combined to create personalized recommendation models, utilizing data from restaurants and the Indonesian food composition list [8]. The system takes into consideration individual variability such as age, sex, physical activity, and medical history to furnish tailored food recommendations, particularly catering to consumers with conditions like diabetes, hypertension, and cardiovascular disease. Through the utilization of a genetic algorithm, the study aims to strike a balance between nutritional adequacy and individual preferences in food selection.

Collaborative filtering suggests meals based on similar user preferences but faces the cold-start problem with new users or items lacking data.

Content-based filtering recommends meals similar to those liked in the past but may limit variety.

Hybrid systems merge collaborative and content-based filtering for more diverse and accurate recommendations.

Within the realm of food recommendation, meal recommendation has emerged as a focal point, concentrating on presenting courses from specific categories for users to enjoy as a complete meal. Previous research efforts have recognized the importance of integrating health and nutrition considerations into meal recommendation systems to encourage healthier eating habits [9]. To address this, researchers have developed various models such as DIETOS, multi-objective optimization technology, and evolutionary approaches, aiming to provide personalized and health-conscious dietary recommendations. The evolution of meal recommendation research has increasingly emphasized health-oriented recommendations, underscoring the necessity of incorporating healthiness criteria into food suggestions.

Existing culinary apps cannot identify accurately available ingredients, leading to ingredient substitutions or missing items in recipes. Food ingredient recognition is achieved through a convolutional neural network model (CNN Model), specifically ResNet-50, trained on a dataset containing 36 classes of vegetables and fruits [10]. Recipe recommendations are made based on the recognized ingredient labels, utilizing the Edamam API to search for suitable recipes.

Clustering algorithms help in organizing and categorizing foods based on their nutritional content, enabling the recommendation system to suggest food options that align with the user's dietary requirements and preferences effectively.

These studies highlight the significance of personalized food recommendations aligned with user preferences and nutritional requirements, employing clustering techniques and data mining approaches.[12]

2.2 Recommendation Systems Based on Deep Learning Algorithms

Lately, there has been a surge in interest in deep learning, primarily because of its capacity to adjust to changing data dynamics and its remarkable versatility in identifying complex patterns and connections in user-item interactions and content data [14,15]. This method satisfactorily addresses the needs of contemporary recommendation algorithms in the era of large-scale data.

Deep learning automates feature extraction by using neural networks to identify and extract relevant patterns from input data without manual intervention. Through multiple layers, each processing data to find increasingly complex features, deep learning models discover intricate patterns and relationships. This capability allows them to handle high-dimensional data effectively, improving performance in tasks like recommendation systems.

Deep learning addresses challenges like the cold-start problem, where new users or objects lack interaction history. By generalizing from existing data, it aids in recommendations where traditional methods struggle. This makes deep learning crucial in recommendation systems, leading to more precise, flexible, and data-driven suggestions. Its ability to navigate complex data landscapes enhances user engagement and experiences. Moreover, it enables the exploration of advanced concepts like reinforcement-learning-based recommendation systems.

2.3 Reinforcement-Learning (RL)-Based Recommendation Systems

Deep learning has established itself as a potent tool for unraveling insights from intricate datasets. However, the pursuit of innovation in recommendation systems extends beyond conventional methods, delving into the realm of reinforcement learning. This method introduces a distinct framework comprising states, actions, environments, agents, and rewards, giving a fresh perspective on optimizing recommendations. Within reinforcement-learning-based recommendation systems, user interactions with the recommendation platform are conceptualized as sequential decision-making processes. Here, the "state" represents the user's prevailing context, "actions" denote the recommendations made to the user, and the "environment" encompasses the recommendation system itself. An "agent" is responsible for generating recommendations based on the user's state and past interactions, while "rewards" serve as feedback on the quality and relevance of these recommendations. This approach enables recommendation systems to evolve, leveraging user feedback and interactions to continually enhance the pertinence and engagement of suggested content.

Therefore, the studies so far highlight the importance of integrating user preferences and nutritional information into food recommendation systems to enhance user satisfaction and promote adherence to healthier dietary habits. By synthesizing insights from diverse research works, this thesis aims to contribute to our understanding of user satisfaction and adherence within the context of these innovative systems, paving the way for enhanced culinary experiences and improved health outcomes.

3. Methodology

This section explains the different aspects of comprehensive recommendation systems, including the problem context, architectural design, and algorithmic framework. By utilizing advanced techniques like the actor-critic methodology in deep reinforcement learning, we implement the DDPG model to improve how the system achieves rewards and shapes its policies. Our main goal is to cater to actively engaged user demographics more effectively.

3.1. Addressing Key Scenarios

In the world of food recommendation systems, conventional models often find it challenging to keep up with users' changing tastes and dietary preferences. We aim to create a recommendation system that goes beyond just looking at a user's past interactions with food. Instead, we want it to continuously learn and improve its suggestions based on the user's evolving dietary preferences and health objectives. This means that food items similar to a user's favorites are recommended, taking into account their latest culinary preferences, nutritional requirements, and any dietary restrictions they may have.

We plan to utilize deep reinforcement learning, particularly the Deep Deterministic Policy Gradient (DDPG)method, to create a dynamic system. This system will continuously learn from user interactions, improve recommendations based on real-time feedback, and adjust to changing user preferences. By representing recommendations as a Markov Decision Process, our goal is to provide personalized and adaptable food suggestions that optimize user satisfaction and align with their dietary objectives and health goals.

3.2 System Framework

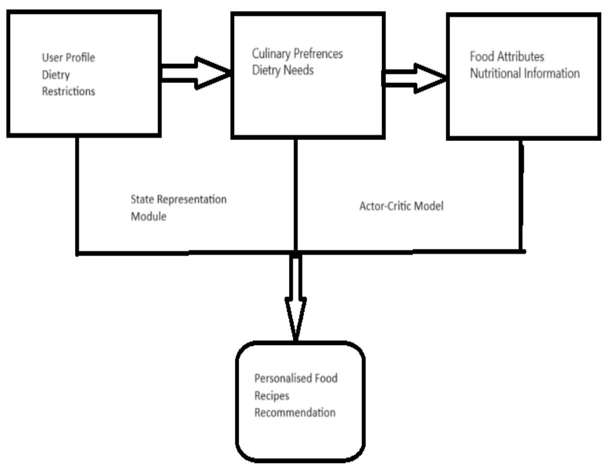

Fig 1. System Framework

We delineate the system framework in Figure 1, illustrating the combination of policies and values intrinsic to the Actor-Critic Model, tailored for the domain of food recommendation. Within this framework, the recommendation system functions based on the Actor-Critic model, wherein the actor makes the policy decisions and the critic evaluates how well they work.

At the heart of this model lies the state representation module, which captures various aspects of the user's culinary preferences and dietary needs. This module includes past interactions, user profiles, contextual cues, and food attributes. Using this comprehensive representation, the actor network creates personalized food recommendations for the user in question. Meanwhile, the Critic network evaluates the actor's decisions by estimating the expected reward or value linked to the action, taking into account the user's current situation.

Moreover, the system integrates user feedback—such as recipe ratings, ingredient selections, and dietary restrictions—as a reward signal, enabling continuous updates to both the actor and the critic. This feedback loop facilitates adaptive learning from user interactions, improving the system's ability to offer tailored food recommendations over time. In food recommendation systems, policy learning (actor) directly links state actions, efficiently making recommendations within extensive action spaces. Value estimation (critic) minimizes update variance, expediting learning and ensuring system stability. Actor-critic methods like DDPG excel in continuous action spaces, balancing exploration-exploitation trade-offs for tailored suggestions. DDPG's sample efficiency and versatility allow effective learning from limited data, making it well-suited for dynamic food recommendation environments.

Policy learning (actor) involves directly acquiring a policy that maps state actions, proving advantageous in systems with large and complex action spaces, such as recommending food items. This method proves more efficient than value-based approaches, allowing the actor to make recommendations according to user preferences and dietary requirements.

Value estimation (critic) plays a crucial role in policy by reducing update variance, making learning faster, and ensuring consistent training. In food recommendation systems, where offering personalized meal suggestions is key, combining actor and critic components enhances the system's ability to adjust to user preferences over time.

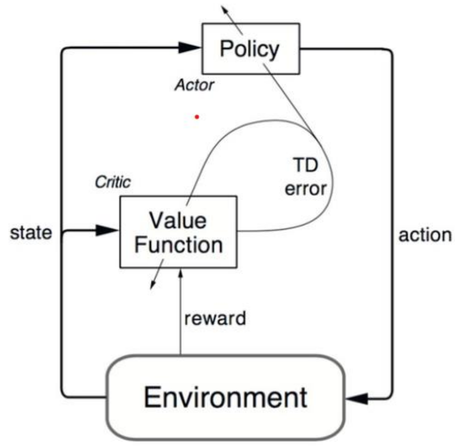

Fig 2 Actor-critic Model, source: [18]

Actor-critic methods like DDPG are particularly suitable for food recommendation systems, handling continuous action spaces and balancing the exploration-exploitation trade-off effectively. DDPG's efficiency in utilizing past experiences allows for learning from limited data, making it perfect for crafting personalized recommendations.

Furthermore, DDPG's versatility adapts to various operational scenarios in food recommendation, allowing training in both real-time user interactions and batch-processing setups. The integration of DDPG into the food recommendation system will be further detailed during implementation, showcasing its effectiveness in handling dynamic environments and providing personalized recommendations. Further details on DDPG integration[18] and its benefits will be elaborated in the implementation phase.

3.3 Implementation

Step 1

- Data Preparation
  The Food.com dataset is sourced to commence the analysis. This dataset comprises a vast collection of recipes, user reviews, and recipe metadata. This initial phase involves preprocessing the dataset, including tasks such as cleaning, organizing, and structuring the data to facilitate subsequent analysis and modeling.
- Embeddings Generation
  This stage focuses on generating embeddings, which capture latent features of recipes and users in a compact space, enhancing recommendation efficiency. Techniques such as matrix factorization may be employed to derive recipe and user embeddings from the recipe-user interaction matrix. These embeddings are then stored for efficient retrieval during the recommendation process.
- Data Partitioning
  The dataset is partitioned into training and testing sets to evaluate the performance of the recommendation model. The dataset is structured to include essential information such as recipe IDs, user IDs, ratings, and other relevant attributes.

Step Two

- State Representation- In food recommendation systems, accurately capturing the present context is vital for guiding both the Actor and Critic networks. This involves blending unique user traits with their past interactions with food recipes. The state representation module combines user characteristics with recipe attributes, adjusting their dimensions using an Adaptive Pooling block. The "concat and flatten" process merges data from various sources into a unified vector [18]. This unified state then serves to create and assess actions, emphasizing the importance of the state representation process in providing personalized recipe recommendations based on user preferences and recipe characteristics.
- Actor and Critic Initialization Our proposed system of food recommendation system, we initialize an Actor and Critic architecture using DDPG. The actor network begins by defining the neural network architecture, which takes the state (including the user's profile, food items, and historical interactions) as input. To ensure diversity, network weights are initialized with small random values, often employing Xavier initialization to maintain reasonable bounds. Both Actor and Critic networks utilize fully connected (dense) layers, the standard for processing high-dimensional input data. The Actor network comprises three linear layers (linear1, linear2, linear3), responsible for processing the state input and outputting the recommended action. Similarly, the Critic network also utilizes three linear layers (linear1, linear2, linear3), but it processes both the state and action together, outputting a single value estimating the Q-value of the state–action pair. In both networks, the input layer receives either the state (in the actor-network) or the state–action pair (in the Critic network), initiating the feature transformation process. Subsequent layers, often referred to as hidden layers, further process these inputs, facilitating the creation of more intricate representations and interactions among features. These hidden layers are essential for capturing the non-linear relationships present in the data.
  In the actor network, the final layer generates the recommended action, while in the Critic network, it produces the estimated Q-value. These outputs play crucial roles in decision-making (Actor) and

evaluation (Critic). The Actor network generates food recommendations, while the Critic network assesses the quality of these recommendations [18].

Algo 1 Actor-critic algorithm Initialisation
The Actor-critic algorithm [16] for a food recommendation system:

1. Define Actor Network-
- Create a neural network architecture for generating food recommendations.
- Include three linear layers followed by LeakyReLU activation and dropout layers for regularization.
- Define the forward method to process input state information and produce food recommendations.

2. Define Critic Network-
- Develop a neural network architecture for evaluating the quality of food recommendations.
- Consists of three linear layers with LeakyReLU activation and dropout layers, concatenating the state and action before processing.
- Implement the forward method to estimate the value of the state-action pair.

3. Model Initialization-
- Set dimensions based on the food recommendation dataset characteristics.
- Initialize Actor and Critic network instances with specified dimensions and dropout probability.
- Set up target networks for stability during training.

4. Optimizer Definition
- Configure optimizers for updating Actor and Critic network parameters during training, using a defined learning rate.

This algorithm outlines the steps to establish and initialize neural networks for generating and evaluating food recommendations within a reinforcement learning-based food recommendation system.

- **Employing DDPG for Training**

The Deep Deterministic Policy Gradient (DDPG) is a model-free algorithm that combines actor-critic strategies to address challenges in continuous action spaces, making it apt for reinforcement learning scenarios [18]. Derived from Deep Q learning, it effectively handles continuous action spaces by incorporating key components. The training process for DDPG, outlined in Algorithm 2, involves several steps aimed at optimizing the model and enhancing its effectiveness.

Algorithm 2: DDPG Training

1. Networks and Parameters Initialization-

- Begin by initializing the actor_model with input dimension, hidden dimension, and output dimension.
- Next initialize the critic_model with input dimension, output dimension, and hidden dimension.
- Assign actor_model to actor_target and critic_model to critic_target.
- Create a replay buffer to store experiences.

2. Setting Exploration Noise, Discount Factor, Tau (Soft Update Parameter)-

- Specify parameters such as exploration noise, discount factor, and tau for soft updates.

3. Training Loop-

- Iterate through episodes from 1 to max_episodes:
- Initialize exploration noise.
- Observe the initial state.
- Iterate through timesteps from 1 to max_timesteps:
    - Generate action using the actor-network with added exploration noise.
    - Execute action and observe the next state, reward, and termination signal.
    - Store experience in the replay buffer.
    - Sample a batch from the replay buffer.
    - Compute target Q-values using the critic_target network.
    - Update the critic network using the sampled batch and target Q-values.
    - Update the actor-network using the sampled batch and critic_model.
    - Soft update the target networks using the tau parameter.
    - Update exploration noise.
    - If the episode is done, break the loop.
- End episode loop.

This algorithm outlines the training process for the Deep Deterministic Policy Gradient (DDPG) algorithm, it involves iteratively updating the actor and critic networks based on experiences stored in a replay buffer to optimize the networks for improved performance in continuous action spaces.

Step Three

This section presents the algorithm of the proposed approach to develop and assess a food recommendation system using deep reinforcement learning techniques.

Algo 3: Proposed Food Recommendation System

Input: recipes_data, user_preferences

Output: evaluation_criteria, tailored_recommendations, Best-N_recommendations

Procedure:

1: Data loading and preprocessing:

- Load the dataset, recipe data, and user preferences.

- Prepare recipe-user interaction matrix.

2: Partition data into training and testing sets:

- divide users into training and testing groups based on their recipe interactions.

3: Data preparation for deep reinforcement learning:

- Create data loaders for training and testing sets.

4: Reinforcement learning models Initialization:

- Initialize actor and critic models for making and evaluating recommendations.

- Set up target networks for stability.

- Initialize replay buffer to store experiences.

5: State representation Definition:

- Define functions to represent the state, including user preferences and recipe features.

6: Model Training:

- For each training episode:

  - For each batch of data:

    - Compute the state representation.

    - Use the actor model to suggest recipes.

    - Calculate rewards based on user interactions with recommended recipes.

    - Update the replay buffer with experiences.

    - If the replay buffer is sufficiently filled, update the actor and critic models.

7: Model Testing:

- For each batch of test data:

  - Compute the state representation.

  - Generate recommendations using the trained actor model.

  - Evaluate the recommendations against the actual user interactions.

8: Compute evaluation metrics:

- Calculate metrics such as accuracy, precision, and recall to assess the performance of the models.

9: User-specific recommendations generation:

- Select a specific user from the test set.

- Generate personalized recommendations for the chosen user.

10: Recommendation Analysis using cosine similarity:

- Compute cosine similarity between user-specific recommendations.

11: Output Return:

- Provide evaluation criteria, tailored recommendations for the selected user, and top-N best recommendations for general users.

4. Results

- Performance Measures-
  In recommendation systems, the evaluation and optimization of performance typically rely on three key metrics: Precision, Recall, and the F1 Score [18]. These metrics ensure that the systems deliver relevant and comprehensive results to users.

1. Precision@k measures the accuracy of recommendations by determining the proportion of relevant items among the top-k recommendations provided by the system. It gauges how well the recommendations align with the user's preferences.
2. Recall@k assesses the system's ability to capture all relevant items within the recommended list by calculating the percentage of relevant items included in the top-k recommendations.

3. The F1 Score serves as a balanced measure of model accuracy, combining Precision and Recall into a single metric. It is particularly useful when striking a balance between Precision and Recall is crucial.

The formulas for these metrics are as follows:
- Precision@k = Number of relevant items among the top-k recommendations / Number of recommended items @k
- Recall@k = Number of relevant items among the top-k recommendations / Total number of relevant items
- F1 Score = 2 * Precision * Recall / (Precision + Recall)

These metrics are pivotal in evaluating recommendation systems, allowing researchers and developers to fine-tune algorithms for more accurate and personalized recommendations.

We primarily assessed our algorithm's performance using the acceptance score, reflecting user satisfaction with the generated food recipes [17]. This score, rated on a scale of 1 to 5, with 1 indicating dissatisfaction and 5 indicating satisfaction, averaged at 4. Significantly higher than alternative methods, this result underscores the high satisfaction levels our algorithm achieved

We employed a preference score to gauge how well our algorithm's meal plans matched users' preferences. This score, calculated as the average TOPSIS score ranging from 0 (no similarity) to 1 (perfect similarity), reflects the degree of alignment between the meal plans and user preferences.

5. Future Directions

In our upcoming research, there are opportunities to enhance the effectiveness and flexibility of our algorithm by exploring additional factors and addressing biases and limitations. Here are some avenues for future exploration:

- Developing a user interface to simplify user interaction and feedback is important. The algorithm can benefit immensely from an intuitive and user-friendly interface that makes it simple for users to choose, rate, and modify their meal plans [17]. For every meal selection, it should provide nutritional statistics, user ratings, and other pertinent information in addition to offering comments and feedback. Monitoring the algorithm's overall development and performance would also be helpful.
- Another topic worth investigating is the integration of data from social media and the Internet of Things (IoT) for contextual suggestions. These platforms may collect and combine data from several sources, including social media interactions, Internet of Things devices, and location-based services, to provide a more comprehensive knowledge of user context. Recommender systems can make recommendations that reflect users' true preferences thanks to this contextual information.
- The digital world today is evolving each second towards more multimodal content, including text, images, and videos. This evolution thereby forces recommender systems to expand their capabilities beyond text-only interactions. By looking at visual content, these systems must gain insights into users' preferences.
- With the integration of GPT-based chatbots combining natural language processing (NLP) with recommendation technology, fundamentally changing the way recommendations are presented. By engaging users in natural conversations, these chatbots improve interaction and personalization, enhancing the overall user experience.

6. Conclusion and Discussion

In this paper, we conclude by suggesting how we can overcome the limitations of traditional food recommendation systems by using a dynamic and customized recommendation framework using deep reinforcement learning, particularly the Deep Deterministic Policy Gradient (DDPG) method [18]. The crux of the whole discussion involves how the utilization of advanced reinforcement learning in recommendation systems. By applying the Actor-critic model for deep learning then supported by the DDPG algorithm may excel in handling continuously changing user preferences action spaces and balancing exploration-exploitation trade-offs effectively, making them well-suited for dynamic food recommendation environments. The power of DDPG lies in efficient learning where limited data or no data is available, making customized recommendations for users.
In our research methodology, we propose the usage of the actor-critic technique within the Deep Reinforcement Learning (DRL) algorithm [18] which plays a very crucial role. This approach combines both policy-based and value-based strategies, which has the potential to improve the whole recommendation process substantially. The beauty lies here in the incorporation of the Deep Deterministic Policy Gradient (DDPG) method supports ongoing and constantly changing interactions between users and food recipes. This

results in such a dynamic system that can cater to individual user preferences effectively. Additionally, by using Singular Value Decomposition (SVD) for matrix factorization in Collaborative Filtering (CF), we have achieved further advancements in innovation. By extracting precise embeddings that capture important user and food features, our method improves recommendation outcomes and customization accuracy significantly.
An adaptive state representation obtained by SVD is also included as part of our technique. By adapting to user choices, this dynamic method makes it possible to give suggestions that are not only more accurate but also customized to the tastes and concerns of individual users.

References

1. Ducrot, P.; Méjean, C.; Aroumougame, V.; Ibanez, G.; Allès, B.; Kesse-Guyot, E.; Péneau, S. Meal Planning Is Associated with Food Variety, Diet Quality and Body Weight Status in a Large Sample of French Adults. Int. J. Behav. Nutr. Phys. Act. 2017, 14, 12.[CrossRef] [PubMed]

2. Stefanidis, K.; Tsatsou, D.; Konstantinidis, D.; Gymnopoulos, L.; Daras, P.; Wilson-Barnes, S.; Dimitropoulos, K. PROTEIN AI Advisor: A Knowledge-Based Recommendation Framework Using Expert-Validated Meals for Healthy Diets. Nutrients 2022, 14,4435. [CrossRef] [PubMed]

3. Curll, M.; DiNardo, M.; Noschese, M.; Korytkowski, M.T. Menu Selection, Glycaemic Control and Satisfaction with Standard and Patient-Controlled Consistent Carbohydrate Meal Plans in Hospitalised Patients with Diabetes. BMJ Qual. Saf. 2010, 19, 355.[CrossRef] [PubMed]

4. Lokuge, C.; Ganegoda, G.U. Implementation of a personalized and healthy meal recommender system in aid to achieve user fitness goals. In Proceedings of the 2021 International Research Conference on Smart Computing and Systems Engineering(SCSE), Colombo, Sri Lanka, 16 September 2021; Volume 4, pp. 84–93. [CrossRef]

5. Raciel yera toledo, ahmad a. Alzahrani, and luis martinez , (Member, IEEE)- A Food Recommender System Considering Nutritional Information and User Preferences.

6. Mehrdad Rostami, Kamal Berahmand, Saman Forouzandeh, Sajad Ahmadian, Vahid Farrahi, Mourad Oussalah - A novel healthy food recommendation to user groups based on a deep social community detection approach.

7. G Prabhakar, Vaishnavi Reddy Chintala, Trishala Reddy, T Ruchitha- User-cloud-based ensemble framework for type-2 diabetes prediction with Diet plan suggestion.

8. Kudang Boro, Evy Damayanthi, Nugraha Edhi Suyatma, Karlisa Priandana, Harry Imantho, Bonang Waspadadi Ligar, Annisa Utami - An Intelligent Food Recommendation System for Dine-in Customers with Non-Communicable Diseases History.

9. Agapito G., Calabrese B., Guzzi P. H., Cannataro M., Simeoni M., Care I., Lamprinoudi T., Fuiano G., Pujia A. - DIETOS: a recommender system for adaptive diet monitoring and personalized food.

10. Miguel Simões Rodrigues, Filipe Fidalgo and Ângela Oliveira - RecipeIS Recipe Recommendation System Based on Recognition of Food Ingredients.

11. Mahdi Nsaif Jasim, Ahmed Bahaddin Hamid -Food recommendation system based on nutritional needs of human beings and user preferences

12. M. Phanich, P. Pholkul, and S. Phimoltares, 'Food recommendation system using clustering analysis for diabetic patients,' in Proc. Int. Conf. Inf.Sci. Appl. (ICISA), Apr. 2010, pp. 1–8.

13. N. Suksom, M. Buranarach, Y. M. Thein, T. Supnithi, and P. Netisopakul, 'A knowledge-based framework for the development of personalized food recommender system,' in Proc. 5th Int. Conf. Knowl., Inf. Creativity Support Syst., 2010, pp. 1–4.

14. Zhang, S.; Yao, L.; Sun, A.; Tay, Y. Deep learning based recommender system: A survey and new perspectives. ACM Comput.Surv. (CSUR) 2019, 52, 1–38. [CrossRef]

15. Cheng, H.T.; Koc, L.; Harmsen, J.; Shaked, T.; Chandra, T.; Aradhye, H.; Anderson, G.; Corrado, G.; Chai, W.; Ispir, M.; et al.- Wide & deep learning for recommender systems. In Proceedings of the 1st Workshop on Deep Learning for Recommender Systems, Boston, MA, USA, 15 September 2016; pp. 7–10.

16. LI Shuang, YAN Yanghui, REN Ju, ZHOU Yuezhi, ZHANG Yaoxue. A Sample-Efficient Actor-Critic Algorithm for Recommendation Diversification[J]. Chinese Journal of Electronics, 2020, 29(1): 89-96. doi: 10.1049/cje.2019.10.004

17. Maryam Amiri, Fatemeh Sarani Rad and Juan Li - Delighting Palates with AI: Reinforcement Learning's Triumph in Crafting Personalized Meal Plans with High User Acceptance

18. Sony Peng 1, Sophort Siet 1, Sadriddinov Ilkhomjon 1, Dae-Young Kim 2,* and Doo-Soon Park 2,*- Combining Deep Reinforcement Learning with Collaborative Filtering for Movie Recommendations