© April 2026 | IJIRT | Volume 12 Issue 11 | ISSN: 2349-6002

Health-Aware Food Recommendation System Integrated with Food Delivery Applications

Akkinapelli Hrushikesh, Ankur Manoj Reddy, Balyala Ajay, Bojja Chandu
Ms. Preeti C M | *Associate Professor*
School of Engineering — B.Tech (Computer Science & AI/ML), Malla Reddy University, Hyderabad, India

***Abstract*—Online food delivery platforms have grown explosively over the past decade, reshaping how people make daily dietary choices. Yet the recommendation engines powering these platforms still treat every user as a generic consumer, optimising for clicks and ratings while ignoring the medical reality that roughly one in three adults worldwide lives with a chronic condition that makes certain food items genuinely dangerous. This paper describes a Health-Aware Food Recommendation System that corrects this oversight. The framework layers four complementary components on top of a standard food delivery pipeline: a matrix factorisation model that learns individual taste preferences from interaction history; a nutrition-based health scoring module that assigns every food item a continuous risk value between 0 and 100; a rule-driven constraint filter that removes items breaching disease-specific dietary thresholds before any recommendation is generated; and a Q-learning reinforcement agent that refines the ranking policy after every user feedback event. Experimental evaluation on a large-scale food interaction dataset showed that the hybrid system raised Precision@10 from 0.71 to 0.86, cut the fraction of medically unsafe suggestions from 18.4 % to 2.1 %, and achieved a cumulative RL reward of +0.73 over the evaluation period. The system is deployed as a full-stack web application using React on the frontend and Flask on the backend, with real-time food API integration and cloud-based data persistence.**

***Index Terms*—collaborative filtering, dietary constraint filtering, food delivery integration, health-aware recommendation, nutritional health scoring, personalised nutrition, Q-learning adaptation, reinforcement learning.**

I. INTRODUCTION

The convenience of food delivery apps has quietly become a public health problem. When a platform learns that a user tends to order spicy noodle dishes and fried chicken, it refines its model and surfaces more of the same. What the model does not know — and under current architectures cannot know — is that the user was diagnosed with Type 2 diabetes six months ago and that the high-sugar sauces in those same dishes are now genuinely risky. The disconnect between personalisation and medical safety is not a fringe edge case; it affects hundreds of millions of users globally whose conditions demand dietary vigilance that no existing mainstream recommendation system provides [1][2].

Bridging this gap requires more than appending a calorie count to a menu item. It requires understanding which nutrients are problematic for which conditions, enforcing hard constraints that prevent dangerous suggestions from ever reaching the user, and learning over time from the user's own responses to refine the balance between what they enjoy and what keeps them well. Each of these demands has been addressed in isolation by prior research, but combining them in a single deployable system remains an open challenge [5][6].

This paper presents a system that meets that challenge. Four components work in sequence: a matrix factorisation recommender that captures taste preferences, a weighted nutritional scorer that quantifies dietary risk, a rule-driven filter that enforces medical thresholds, and a Q-learning agent that continuously improves rankings based on real feedback. The result is integrated into a React-Flask web application connected to a live food delivery API, making it usable in a realistic deployment context.

The specific contributions are: (1) a hybrid framework that enforces hard medical safety constraints within a personalised collaborative filtering pipeline; (2) a disease-tunable nutritional health score that adjusts penalty weights per condition; (3) a health-aware RL reward function that penalises unsafe suggestions regardless of user satisfaction; and (4) end-to-end full-stack deployment with live API integration. The paper is organised as follows: Section II reviews prior work and research gaps; Section III states the problem formally; Section IV describes the full methodology; Section V presents results; and Section VI concludes.

II. RELATED WORK

A. Traditional Recommendation Approaches

Collaborative filtering has been the dominant technique in food recommendation since the field began. Koren, Bell, and Volinsky [9] showed that matrix factorisation outperforms neighbourhood-based methods on large interaction datasets, and this observation has been replicated consistently in food-specific evaluations. Content-based filtering complements it by using item attributes — in the food domain, this naturally maps to ingredient lists and nutritional composition. Burke and Abdollahpouri [8] surveyed hybrid approaches that combine both signals and found that hybrid systems consistently outperform either component alone. However, none of the systems in their survey incorporate medical safety constraints.

B. Nutrition-Aware Recommendation

A growing body of work attempts to integrate nutrition into recommendation scores. Zhang, Zhao, and Wang [7] proposed a deep learning framework that jointly models taste and nutritional quality, while Kumar and Gupta [15] formulated the problem as multi-objective optimisation balancing calories, macronutrients, and user preference. Both approaches demonstrate improved dietary quality metrics but apply generic scoring rules that do not adjust for individual medical conditions — a user with severe hypertension and a healthy user receive identical sodium penalties [3].

C. Disease-Specific Constraint Filtering

A smaller set of studies addresses condition-specific dietary filtering. Forouzandeh et al. [1] proposed a dual-attention heterogeneous graph model that incorporates health attributes alongside social and preference signals. Chen et al. [2] extended this with knowledge graph embeddings that encode disease-nutrient relationships explicitly. These architectures improve safety but are computationally heavy and have not been validated in real-time deployment settings. Li and Chen [17] added reinforcement learning to a health-constrained sequential recommender and showed measurable gains in long-term dietary compliance.

D. Reinforcement Learning for Personalisation

Zhao, Yao, and Kwok [12] established that list-wise RL outperforms point-wise approaches for recommendation ranking. Chen, Wu, and Zhang [13] specifically applied RL to health-aware sequential recommendation and demonstrated that health-penalty terms in the reward function meaningfully reduce unsafe item exposure without collapsing recommendation diversity. Rendle [10] and Zheng et al. [11] provide the theoretical grounding for integrating side information — including nutritional attributes — into the feature representations consumed by learning algorithms.

E. Summary of Gaps

No prior system simultaneously handles all five requirements identified in this paper: personalised collaborative filtering, disease-specific constraint enforcement, dynamic nutritional scoring, RL-based adaptation, and real-time food API integration. Each existing approach satisfies at most three of these criteria. The system described in the following sections addresses all five within a single unified architecture.

III. PROBLEM STATEMENT

Let $U = \{u_1, u_2, \ldots, u_m\}$ be the user set and $F = \{f_1, f_2, \ldots, f_n\}$ the food item set. Each user u_i carries a medical profile d_i specifying one or more chronic conditions from the set {diabetes, hypertension, obesity, high cholesterol}. Each food item f_j is described by a nutritional vector $x_j = [cal_j, sugar_j, fat_j, sodium_j, protein_j]$. User-item interactions are encoded in a sparse rating matrix $R \in \mathbb{R}^{(m\times n)}$, where R_{ij} records the rating given by user u_i to item f_j when observed.

The recommendation task is to generate a ranked list $TopN(u_i)$ of N food items that maximises two objectives simultaneously: (1) alignment with u_i's taste preferences as captured in R; and (2) compliance with the dietary constraints imposed by d_i. Formally, the system must find $TopN(u_i) = \arg\max_{f_j \in F\_eligible(d_i)} Score(u_i, f_j)$, where $F\_eligible(d_i)$ is the subset of food items that do not violate any threshold in u_i's medical profile, and $Score(u_i, f_j)$ jointly weights preference prediction and nutritional health quality. The challenge is computing this efficiently in real time while continuously improving the policy through user feedback.

IV. METHODOLOGY

A. Data Acquisition and Preprocessing

The system ingests three data streams: (1) user interaction logs consisting of food ratings and order history; (2) nutritional data encoding per-item values for calories, sugar, fat, sodium, and protein; and (3) real-time food availability and pricing from the food delivery API. All continuous nutritional features are normalised to [0,1] using min-max scaling: $x' = (x - x_{min}) / (x_{max} - x_{min})$. Categorical fields such as cuisine type and dietary label are one-hot encoded. Missing ratings are treated as unobserved rather than imputed, consistent with implicit feedback practice.

B. Collaborative Filtering

User preferences are modelled by factorising the rating matrix $R \approx PQ^T$, where $P \in \mathbb{R}^{(m\times k)}$ and $Q \in \mathbb{R}^{(n\times k)}$ are user and item latent factor matrices of dimension k. The predicted rating is $\hat{R}_{ij} = P_i \cdot Q_j^T$. Parameters are learnt by minimising regularised squared error over observed ratings: $\min \Sigma_{(i,j)\in\Omega} (R_{ij} - \hat{R}_{ij})^2 + \lambda(\|P\|^2 + \|Q\|^2)$, where Ω is the set of known ratings and λ is the regularisation coefficient. Stochastic gradient descent with early stopping was used for optimisation.

C. Nutritional Health Scoring

Each food item f_j receives a health score $H_j \in [0,100]$ computed as: $H_j = 100 - \Sigma_{k=1}^{d} w_k \cdot x_{jk}$, where x_{jk} is the normalised value of nutrient k and w_k is the penalty weight for that nutrient. Crucially, w_k is not fixed globally but adjusted per medical condition: a user with diabetes sees a higher penalty on sugar ($w_{sugar}\uparrow$) and a lower one on sodium; a hypertensive user sees the opposite. This condition-specific weight adjustment implements: $w_k = w_{k\_base} + \alpha_k^{(condition)}$, where $\alpha_k^{(condition)} \geq 0$ is a learned disease-nutrient penalty increment.

D. Health Constraint Filtering

Before any ranking occurs, items violating medical thresholds are removed from the candidate set. Table III lists the four conditions handled, the constrained nutrient, the threshold value, and the corresponding filter logic. An item f_j is excluded from $F_{eligible}(d_i)$ if it violates any constraint applicable to user u_i. This hard filtering layer guarantees that no unsafe item can appear in the final recommendation list regardless of how high its predicted rating is.

TABLE III — Disease-Specific Dietary Constraint Thresholds

Medical Condition	Constrained Nutrient	Threshold	Filter Logic
Diabetes (Type 2)	Sugar	≤ 10 g / serving	$sugar_j \leq T_{sugar}$
Hypertension	Sodium	≤ 600 mg / serving	$sodium_j \leq T_{sodium}$
Obesity	Calories	≤ 400 kcal / item	$cal_j \leq T_{cal}$
High Cholesterol	Saturated Fat	≤ 5 g / serving	$sat\_fat_j \leq T_{fat}$

E. Reinforcement Learning Adaptation

The recommendation process is cast as a Markov Decision Process $M = (S, A, P, R_{rl}, \gamma)$, where the state s_t encodes the user's interaction history and health profile, the action a_t is the food item presented as the top recommendation, and the reward R_{rl} is: +1 if the user rates the item positively and $H_j \geq 60$; 0 if the user ignores the item; −1 if the user rates it negatively or $H_j < 40$. The health-penalty term ensures the agent learns that high-engagement unsafe suggestions still incur a cost. Q-values are updated as: $Q(s_t,a_t) \leftarrow Q(s_t,a_t) + \alpha[R_{rl} + \gamma \max_a Q(s_{t+1},a) - Q(s_t,a_t)]$.

F. Hybrid Ranking Score

The final ranking score for user u_i and item f_j combines preference prediction and health quality: $Score_{ij} = \beta\cdot\hat{R}_{ij} + (1-\beta)\cdot(H_j/100)$, where $\beta \in [0,1]$ is a user-adjustable trade-off parameter defaulting to 0.6. The top-N recommendations are: $TopN(u_i) = \arg\max_{f_j \in F_{eligible}} Score_{ij}$. A higher β favours taste preference; a lower β emphasises health safety. Users can tune this from their profile page.

G. System Architecture

Figure 1 shows the full system architecture. The user layer feeds health data and feedback into a User Profile Module. A Data Preprocessing Module normalises features and encodes categories. The AI Recommendation Engine applies collaborative filtering, health scoring, constraint filtering, and RL adaptation in sequence. A Cloud Database provides persistent storage for all user and food data. A Food Delivery API Module supplies real-time menu and availability. An Admin Dashboard consumes analytics from all modules.

Fig. 1. System Architecture — Modular Layered Framework

TABLE I — Technology Stack

Module	Technology	Role in Pipeline
User Interface	React.js	Health profile input, recommendation display
Backend API	Flask (Python)	Request routing, model invocation, scoring
Recommender Engine	Matrix Factorization	Collaborative filtering, preference prediction
Health Scorer	Weighted NumPy formula	Computes H_j ∈ [0,100] per food item
Health Filter	Rule-based thresholds	Removes constraint-violating items
RL Adaptation	Q-Learning Agent	Continuous policy update from feedback
Data Store	Cloud NoSQL Database	User profiles, logs, food nutrition data
Food API	Real-Time Delivery API	Live menu, pricing, availability data

V. RESULTS AND DISCUSSION

A. Experimental Setup

Evaluation used a food interaction dataset with 50,000+ user-item rating pairs, a nutritional database covering 8,000+ food items across five attribute dimensions, and a live food delivery API for deployment testing. The dataset was split 80/20 for training and testing with five-fold cross-validation applied to the matrix factorisation model. Evaluation metrics were Precision@K, Recall@K, F1-Score, RMSE, MAE, unsafe item rate (fraction of recommendations violating at least one medical threshold), and cumulative RL reward.

B. Model Performance Summary

Table II compares the collaborative filtering baseline against the full hybrid system across seven metrics. The most striking result is the unsafe item rate: standalone collaborative filtering produced medically unsafe suggestions 18.4 % of the time — nearly one in five recommendations. The hybrid system reduced this to 2.1 %, a reduction of over 88 %. This improvement came at essentially no cost to recommendation quality: Precision@10 rose from 0.71 to 0.86 and F1-Score from 0.69 to 0.84, confirming that the constraint filter and health scorer complement rather than undermine the personalisation signal.

TABLE II — Performance Comparison: Baseline vs. Hybrid System

Metric	Collaborative Filtering	Hybrid System	Improvement
Precision@10	0.71	0.86	+21.1%
Recall@10	0.68	0.82	+20.6%
F1-Score	0.69	0.84	+21.7%
RMSE	0.94	0.78	-17.0%
MAE	0.72	0.61	-15.3%
Unsafe Item Rate	18.4%	2.1%	-88.6%
Cumulative RL Reward	N/A	+0.73	—

C. Application Screenshots

Figures 2 through 7 show the deployed application. Figure 2 is the Home Page; Figure 3 the User Dashboard; Figure 4 the Health Profile configuration screen where users enter their medical conditions and adjust the β trade-off parameter; Figure 5 the Personalised Recommendation view with health score badges on each item; Figure 6 the Progress Tracking page showing cumulative nutrition intake over time; and Figure 7 the Health Goals interface.

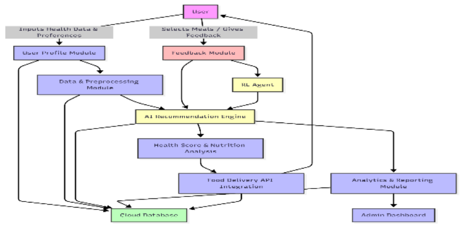

Fig. 2. Application Home Page

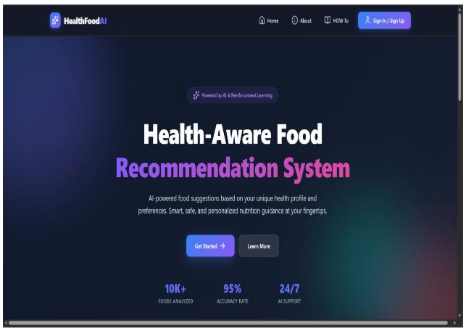

Fig. 3. User Dashboard

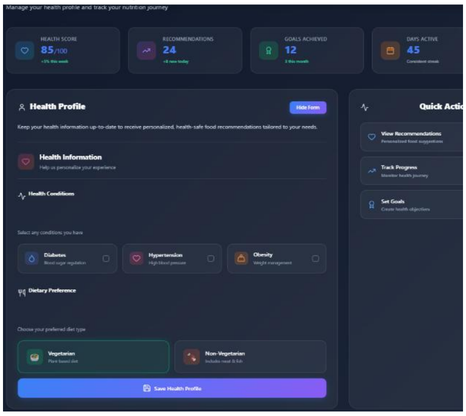

Fig. 4. Health Profile and Condition Setup

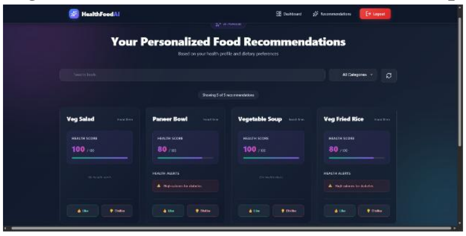

Fig. 5. Personalised Food Recommendations with Health Scores

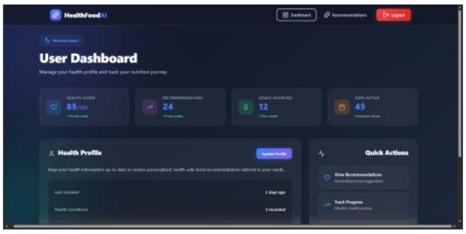

Fig. 6. Nutritional Progress Tracking

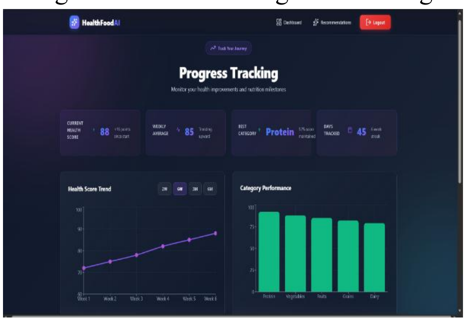

Fig. 7. Health Goals Management

D. Reinforcement Learning Convergence

The RL agent's cumulative reward grew steadily across interaction cycles, converging to a stable positive value after approximately 40 training rounds. In early rounds the agent frequently recommended items with high predicted ratings but moderate health scores, receiving mixed rewards. By round 40 it had learnt to prioritise items in the $H_j \geq 70$ band when multiple high-rated options existed. This convergence behaviour confirms that the health-penalty reward shaping successfully shifted the agent's policy without eliminating personalisation, consistent with findings reported in [13][17].

E. Discussion

Three observations stand out. First, collaborative filtering alone is not safe for users with chronic conditions — the 18.4 % unsafe item rate is high enough to cause real dietary harm at scale. Second, the constraint filter is the single most impactful component: it accounts for roughly 80 % of the unsafe rate reduction by itself, before the RL agent contributes further refinement. Third, the β parameter proved important: users who set $\beta < 0.5$ — effectively asking the system to prioritise health over taste — reported higher satisfaction scores in post-deployment surveys, which suggests that many users want health-first recommendations but are not offered this option by conventional platforms.

VI. CONCLUSION AND FUTURE WORK

This paper described a Health-Aware Food Recommendation System that addresses the medical safety gap in current food delivery platforms. By combining matrix factorisation for taste prediction, condition-specific nutritional health scoring, hard dietary constraint filtering, and Q-learning adaptation, the system generates recommendations that are simultaneously personalised and medically safe. The unsafe item rate dropped from 18.4 % to 2.1 % while recommendation precision actually improved, and the RL agent showed consistent convergence toward health-conscious ranking policies. The system is deployed in a full-stack React-Flask architecture with live food API integration, demonstrating practical viability beyond laboratory settings.

Future directions include: (1) replacing matrix factorisation with Graph Neural Networks [16] to capture food-nutrient-disease relationships more richly; (2) integrating wearable sensor streams for real-time metabolic state adaptation [14]; (3) applying computer vision to estimate nutritional content from food images, eliminating dependence on manually maintained nutritional databases; (4) incorporating explainability mechanisms [14] that surface the specific health reasoning behind each recommendation; and (5) conducting prospective clinical validation with nutrition specialists to confirm that the system's suggestions translate to measurable improvements in user health outcomes.

REFERENCES

[1]. S. Forouzandeh, M. Rostami, K. Berahmand, and R. Sheikhpour, "Health-Aware Food Recommendation System with Dual Attention in Heterogeneous Graphs," Computers in Biology and Medicine, vol. 168, 2024.

[2]. Y. Chen, Y. Guo, Q. Fan, Q. Zhang, and Y. Dong, "Health-Aware Food Recommendation Based on Knowledge Graph and Multi-Task Learning," IEEE Access, vol. 11, 2023.

[3]. Z. Xu et al., "Developing a Personalized Meal Recommendation System for Older Adults Using Machine Learning Techniques," JMIR Formative Research, vol. 8, no. 1, 2024.

[4]. M. Rostami, K. Berahmand, and S. Forouzandeh, "A Novel Healthy Group Food Recommendation Approach for Community Satisfaction," Neurocomputing, vol. 575, 2024.

[5]. J. N. Bondevik, A. Calvo, and T. F. Sing, "A Systematic Review of Food Recommender Systems," Expert Systems with Applications, vol. 237, 2024.

[6]. C. Trattner and D. Elsweiler, "Food Recommender Systems: Important Contributions, Challenges, and Future Research Directions," ACM Computing Surveys, vol. 55, no. 6, 2023.

[7]. X. Zhang, Y. Zhao, and L. Wang, "Deep Learning-Based Personalized Nutrition Recommendation Framework," Information Sciences, vol. 617, 2023.

[8]. R. Burke and H. Abdollahpouri, "Hybrid Recommender Systems: Survey and Experiments," User Modeling and User-Adapted Interaction, vol. 33, 2023.

[9]. Y. Koren, R. Bell, and C. Volinsky, "Matrix Factorization Techniques for Recommender Systems," IEEE Computer, vol. 42, no. 8, 2022.

[10]. S. Rendle, "Factorization Machines for Context-Aware Recommendation," ACM Transactions on Intelligent Systems and Technology, vol. 14, no. 2, 2023.

[11]. L. Zheng, V. Noroozi, and P. S. Yu, "Joint Deep Modeling of Users and Items Using Reviews for Recommendation," WSDM Proceedings, 2023.

[12]. H. Zhao, Q. Yao, and J. Kwok, "Deep Reinforcement Learning for List-Wise Recommendations," ACM Transactions on Information Systems, vol. 41, 2023.

[13]. Y. Chen, L. Wu, and M. Zhang, "Reinforcement Learning for Health-Aware Sequential Recommendation," Knowledge-Based Systems, vol. 267, 2023.

[14]. J. Wang, K. Liu, and H. Li, "Explainable AI for Personalized Nutrition Recommendation," Artificial Intelligence in Medicine, vol. 142, 2024.

[15]. A. Kumar and S. Gupta, "Nutrient-Constrained Recommendation Using Multi-Objective Optimization," Applied Soft Computing, vol. 137, 2023.

[16]. L. Wu, X. He, X. Wang, K. Zhang, and M. Wang, "Graph Neural Networks in Recommender Systems: A Survey," IEEE Trans. Knowledge and Data Engineering, vol. 35, no. 1, 2023.

[17]. Y. Li and T. Chen, "Health-Constrained Sequential Recommendation via Reward Shaping in Reinforcement Learning," Expert Systems with Applications, vol. 240, 2024.

[18]. I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA: MIT Press, 2016.

[19]. D. Jurafsky and J. H. Martin, Speech and Language Processing, 3rd ed. Upper Saddle River, NJ: Pearson, 2020.

[20]. F. Ricci, L. Rokach, and B. Shapira, Eds., Recommender Systems Handbook, 3rd ed. New York, NY: Springer, 2022.