NutriGuardian AI: An Agentic Generative System for Sustainable Nutrition

Neha Ravi Kumar
Department of Artificial Intelligence and Data Science,
East West Institute of Technology,
Bengaluru, India
Email: nehardxb23@gmail.com

Prof. Akshatha K S
Dept. of Artificial Intelligence and Data Science
East West Institute of Technology,
Bengaluru, India
Email: akshathaks@ewit.edu.in

Shreya G
Department of Artificial Intelligence and Data Science,
East West Institute of Technology,
Bengaluru, India
Email: shreyayash.1207@gmail.com

Kaveri M H
Department of Artificial Intelligence and Data Science,
East West Institute of Technology,
Bengaluru, India
Email: kaverimh25@gmailcom

Abstract—Personalized nutrition requires more than a generic calorie chart because dietary requirements depend on body measurements, lifestyle, activity, goals, preferences, affordability, and health conditions. Existing diet applications frequently separate meal planning, health-risk assessment, progress tracking, and conversational assistance, forcing users to interpret information manually. This paper presents NutriGuardian AI, an integrated web-based system that combines machine learning, Generative AI, and an agentic feedback mechanism for sustainable nutrition and preventive healthcare. The system provides secure user registration and profiles, health-risk prediction for diabetes, heart disease, hypertension, and obesity, personalized meal-plan generation, AI-based nutrition conversation, progress monitoring, and sustainability-oriented insights. A React and Tailwind CSS frontend communicates with a FastAPI/Python backend, while MySQL provzdes persistent storage. Disease models are developed using machine-learning algorithms and serialized for API-based inference. The Generative AI layer uses contextual user information to create meal guidance and explanations, while the agentic layer observes progress and can update recommendation context when user state changes. The system is designed around Sustainable Development Goal 2 (Zero Hunger) and SDG 3 (Good Health and Well-being), with emphasis on accessible, affordable, region-friendly, and preventive nutrition support.
The proposed architecture connects prediction, generation, tracking, and adaptation into a single workflow rather than treating them as independent features. The paper describes the system architecture, methodology, database design, implementation, evaluation strategy, limitations, and future enhancements.

Index Terms—Personalized Nutrition, Preventive Healthcare, Machine Learning, Generative AI, Agentic AI, Health Risk Prediction, Diet Planning, FastAPI, React, MySQL, Sustainability, SDG 2, SDG 3.

I. INTRODUCTION

Nutrition is closely associated with physical well-being, disease prevention, energy balance, and quality of life. However, a suitable diet is not identical for every individual. Age, gender, height, weight, activity level, food preference, health status, target weight, lifestyle, and local food availability can all influence nutritional requirements. Traditional diet charts generally

provide fixed recommendations and may not respond when a user's health profile or progress changes. Similarly, basic calorie-tracking applications may record information without combining it with predictive health insights and adaptive meal planning.

Recent developments in artificial intelligence have enabled nutrition systems to move from fixed rule-based recommendations toward personalized and data-driven approaches. Machine-learning models can identify patterns in structured health data, while Generative AI can produce natural-language explanations and flexible meal suggestions. Prior work has explored personalized nutrition recommendations, AI-generated recipes, adaptive diet planners, explainable recommendations, and health-aware food recommendations. The supplied base literature also emphasizes that generated nutritional content should be validated because language models can produce plausible but inaccurate information. [1]–[3]

NutriGuardian AI is designed as a unified preventive nutrition platform. Its objective is not merely to recommend meals, but to connect to a user's profile, predicted health risks, dietary goals, generated meal plans, progress records, and feedback. The project also introduces an agentic layer intended to observe changes over time and use them as context for future recommendations. This creates a feedback loop instead of a one-time recommendation process.

The major contributions of the proposed system are: 1) an integrated architecture combining ML prediction, Generative AI, and agentic adaptation; 2) disease-risk prediction for diabetes, heart disease, hypertension, and obesity; 3) personalized meal planning and AI nutrition conversation; 4) secure authentication and longitudinal progress tracking; 5) sustainability-oriented nutrition insights; and 6) an accessible web interface designed to support preventive and affordable nutrition decisions.

II. MOTIVATION AND PROBLEM STATEMENT

Many existing nutrition solutions address only one part of the user's problem. A calorie counter may monitor energy intake but not estimate disease risk. A recipe application may suggest meals without considering the user's predicted health condition. A static diet chart may be appropriate for one point in time but become less relevant after changes in weight, activity, or goals. Users may therefore need several independent tools and must manually combine their results.

The problem addressed by NutriGuardian AI is the design of a software platform that can transform health and lifestyle information into personalized, explainable, and adaptive nutrition guidance. The platform should maintain a persistent user context, estimate selected health risks, generate suitable meal guidance, record progress, and use feedback to improve subsequent recommendations.

The project is also motivated by accessibility and sustainability. Nutrition recommendations should consider practical constraints such as commonly available foods, dietary preference, approximate cost, and regional suitability instead of assuming that every user can follow expensive or specialized diets. The system therefore treats affordability and

sustainability as supporting dimensions of personalized nutrition.

III. RELATED WORK

A. Personalized Nutrition Recommendation

Machine-learning-based personalized nutrition systems use food and user attributes to classify foods or generate recommendations. Samy et al. considered nutritional variables including calories, fat, cholesterol, sodium, fiber, protein, vitamins, calcium, potassium, and water, and investigated clustering, classification, and explainability for personalized recommendation. [3] Such work supports the use of structured nutritional attributes when designing an intelligent recommendation layer.

B. AI-Generated Recipes

Tanwar et al. studied quantitative evaluation of AI-generated recipes for health recommender systems. Their work highlights the difference that may exist between generated and human-curated recipes and motivates post-generation checking of ingredients and nutrient information. [1] This is important for NutriGuardian AI because Generative AI is used to produce flexible meal guidance rather than only retrieving fixed recipes.

C. Adaptive Diet Planning

The Smart Diet Planner work describes a web-based approach combining personalization, meal planning, tracking, dashboards, and adaptive substitutions. It also discusses energy estimation and feedback-driven recommendations. [2] NutriGuardian AI follows this general adaptive direction but connects it with disease-risk prediction, an AI conversational layer, and an agentic monitoring mechanism.

D. Health-Aware Recommendation

Earlier health-aware food recommendation research has investigated the use of health information and constraints during food selection. Other studies have examined hybrid recommendations, recipe knowledge graphs, explainable recommendations, and generative approaches. [13] These studies demonstrate the importance of constraints, explainability, nutritional information, and user context. The proposed project combines these ideas into a single software workflow.

IV. OBJECTIVES

The primary objective is to develop NutriGuardian AI as an intelligent nutrition and preventive-health support platform. The specific objectives are: (1) collect and securely manage user profile and health information; (2) estimate selected health risks using trained ML classification models; (3) generate personalized meal plans according to user characteristics, goals, and constraints; (4) provide conversational AI assistance and explanations; (5) track weight, calories, meals, and other progress indicators; (6) use an agentic feedback mechanism for adaptive recommendations; (7) provide sustainability-oriented information; and (8) maintain an architecture that can be extended with additional diseases, foods, languages, and data sources.

V. PROPOSED SYSTEM

NutriGuardian AI is implemented as a modular web application. The system begins with registration and authentication. After logging in, the user can create or update a profile containing demographic, body, lifestyle, dietary, and goal-related information. This information becomes a

shared context for the downstream prediction and nutrition modules.

The disease prediction component accepts structured health parameters and produces risk outputs for supported conditions. In the current project scope, the endpoints cover diabetes, heart disease, hypertension, and obesity. The predictions are intended as decision-support indicators and not as clinical diagnoses. The result can be presented together with the relevant risk category and stored in prediction history.

The Smart Diet Planner uses user context and health information to generate a structured nutrition plan. The output may contain breakfast, lunch, dinner, snacks, approximate calories and protein information, hydration guidance, food substitutions, and explanations. The Generative AI component is also used through a conversational interface so that users can ask nutrition-related questions and receive context-aware responses.

The progress-tracking component stores longitudinal observations such as weight, calorie information, meals, and dates. Dashboard visualizations convert these records into understandable trends. The sustainability component provides an additional perspective on food choices.

VI. SYSTEM ARCHITECTURE

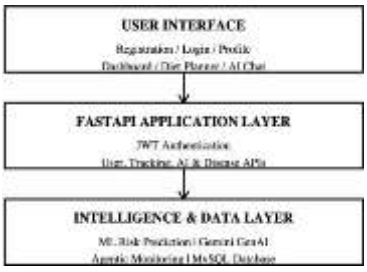


Fig. 1. *High-level architecture of NutriGuardian AI.*

The architecture follows a layered design consisting of the presentation layer, application/API layer, intelligence layer, and data layer. The frontend is developed using React with Tailwind CSS and visualization components such as Chart.js. Framer Motion can be used for interface transitions and interaction feedback. The backend is developed in Python using FastAPI. Machine-learning components are implemented using libraries such as scikit-learn and other model-development frameworks selected for the individual tasks.

The API layer provides authentication, user, tracking, disease prediction, and AI-related services. JWT-based authentication is used to associate protected requests with the authenticated user. The data layer uses MySQL for persistent storage. Model files are loaded by prediction services for inference. The architecture is modular so that a predictive model can be retrained or replaced without redesigning the authentication and tracking layers.

Fig. 1 can be represented as: User Interface → FastAPI Application → Intelligence Services → MySQL / Model Storage. Within the intelligence services, the prediction path contains disease-specific ML models, while the nutrition path contains the Generative AI service.

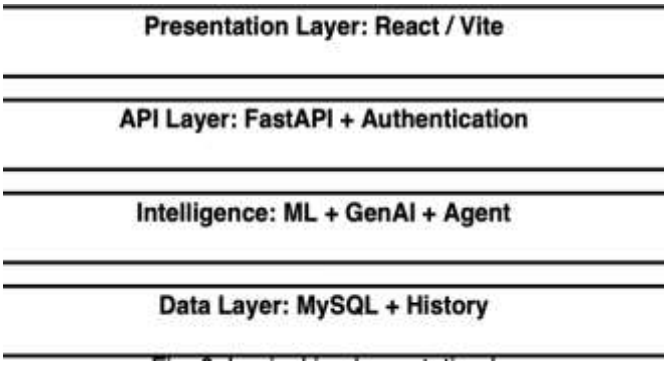


Fig. 2. Logical implementation layers.

VII. METHODOLOGY

A. Data Collection and User Profile

The workflow begins with user registration and profile creation. Relevant attributes include age, gender, height, weight, activity level, dietary preference, health goal, and other constraints required by individual modules. Validation is performed before information is stored. A persistent user identifier links the profile with prediction, tracking, meal, and conversation records.

B. Disease Prediction

The health-risk module follows a supervised classification workflow. Dataset preparation includes loading structured records, handling the feature columns required by each model, training a classifier, evaluating it on held-out data, and serializing the trained model. During inference, the API receives a feature vector, loads the corresponding model, generates the prediction, and returns the result to the frontend.

The current project contains disease-specific model artifacts for diabetes, heart disease, hypertension, and obesity. Encoders are also used where categorical variables require transformation. The diabetes model is based on the PIMA diabetes dataset in the current implementation. The model-development process should report dataset source, feature list, train-test split, preprocessing, algorithm, and measured evaluation metrics in the final experimental record.

C. Personalized Diet Generation

The diet-generation workflow converts user information into a structured prompt. The prompt may contain body measurements, health goals, dietary preference, activity level, predicted risk context, regional food preference, and affordability constraints. The Generative AI service then produces a personalized plan. Structured prompts are important because they constrain the output format and reduce irrelevant responses.

Generated recommendations should be treated as guidance rather than medical prescriptions. Nutrition values generated by a language model may require validation against trusted food-composition information. This follows the concern identified in the AI-generated recipe literature. [1]

D. Agentic Adaptation

The agentic mechanism is represented by a continuous cycle: observe → evaluate → decide → update context → recommend. The system can inspect recent progress, identify changes or deviations from the user's goal, and provide updated information to the recommendation layer. This differs from a static chatbot because the agent uses persistent state and feedback as part of the recommendation process.

E. Tracking and Visualization

Progress records are timestamped and associated with the authenticated user. Weight and calorie information can be summarized across daily, weekly, and monthly periods. Charts help users identify trends rather than relying on a single measurement. Meal history and chat history can also provide context for future personalization.

F. Sustainability

Sustainable nutrition is incorporated as an additional design objective. The system can consider region-friendly food choices and sustainability indicators alongside nutritional requirements. The exact scoring formula and data source should be documented in the final

implementation so that the indicator is reproducible, and its interpretation is clear.



Fig. 3. Personalized recommendation and feedback pipeline.

VIII. TECHNOLOGY STACK AND IMPLEMENTATION

The frontend uses React for component-based interface development and Tailwind CSS for responsive styling. Chart.js supports progress and risk visualizations, while Framer Motion can improve interaction feedback. The backend uses Python and FastAPI to expose REST-style endpoints. Machine-learning functionality can be used using scikit-learn, TensorFlow, XGBoost, or LightGBM depending on the specific model. Computer-vision components such as YOLOv8/EfficientNet/OpenCV are reserved for future or optional food-image analysis rather than being assumed as a completed module.

The database layer uses MySQL. SQLAlchemy can provide object-relational database interaction, while Pydantic/FastAPI schemas can validate API inputs and outputs. Authentication uses JWT and may support Google Login where configured. External Generative AI services such as the Gemini API provide language-model functionality. Google Translate or another translation service can be integrated for multilingual guidance. Deployment can use Docker and cloud platforms such as Render, Railway, or AWS when required.

The modular implementation separates routes, schemas, database models, services, machine-learning models, and AI utilities. This separation supports team development and makes it easier to test individual services. API testing can be performed using standard HTTP clients such as Postman when required, although it is not a core dependency of the application.

IX. DATABASE DESIGN

The database is centered on the user's authentication. A user's table stores profile information. A progress_tracking table stores longitudinal observations. A meal_history table store generated or recorded meal information. A chat_history table stores AI conversations when conversation persistence is enabled. A prediction_history table stores disease prediction outputs and timestamps.

The database design establishes relationships between these tables using the user identifier, allowing information from different modules to be associated with the correct user. This structure enables the system to retrieve previous meals, health predictions, conversations, and progress records whenever required for personalized recommendations.

Table	Important Fields	Purpose
users	id, name, email, password, height, weight, age, gender	Authenticated profile and baseline health information.
progress_tracking	id, user_id, weight, calories, date	Longitudinal progress observations.
meal_history	id, user_id, meal, calories, date	Generated/recorded meal history.
chat_history	id, user_id, question, answer, date	AI nutrition conversation history.
prediction_history	id, user_id, disease, prediction, risk, date	Disease-risk prediction history.

Fig. 4. Database Design

X. FUNCTIONAL MODULES

A. Authentication and User Management: The module provides registration, login, protected access, profile retrieval, and profile updates. User-specific records are linked through the authenticated identifier.

B. Disease Prediction: Four disease-risk workflows are included in the project scope: diabetes, heart disease, hypertension, and obesity. Each workflow maps validated input features to a trained model and returns an interpretable risk result.

C. Smart Diet Planner: The planner uses user context to generate meal schedules and nutritional guidance. It can support substitutions and region-friendly food choices.

D. AI Nutrition Chat: The conversational module allows users to ask questions about meals, nutrition, goals, and recommendations. The user's stored context can be used to make responses more relevant.

E. Progress Dashboard: Weight, calories, meals, and prediction information can be visualized through charts and summary cards. The dashboard supports a longitudinal understanding of progress.

F. Agentic Monitoring: The monitoring layer evaluates stored progress and feedback and supplies updated context to the recommendation system. Its purpose is to reduce dependence on manual re-planning.

G. Sustainability Tracking: Sustainability-oriented indicators provide additional information about food choices and encourage consideration of environmental impact together with health.

XI. SECURITY, PRIVACY, AND RESPONSIBLE AI

Because NutriGuardian AI handles health-related information, security and responsible AI are important design considerations. Authentication must restrict protected resources to the relevant user. Passwords should be stored using secure hashing rather than plaintext. API inputs should be validated, and sensitive configuration values such as API keys and database credentials should be maintained through environment variables.

Health-risk outputs should not be presented as confirmed medical diagnoses. The system should clearly communicate that predictions are model-based estimates and encourage professional consultation when users require medical advice. Generative responses should be treated as supportive guidance. The system should avoid fabricating certainty when the available user data is incomplete.

Privacy should also be considered during storage, logging, testing, and deployment. Only necessary user information should be retained, and access to prediction of history, meals, and conversations should be controlled. Future versions can add stronger consent management, encryption, audit logs, data deletion workflows, and privacy-preserving analytics.

XII. EXPERIMENTAL EVALUATION

Evaluation of the proposed system should be performed at both model and application levels. Disease models can be evaluated using accuracy, precision, recall, F1-score, confusion matrices, and where appropriate ROC-AUC. The evaluation dataset must remain separate from training data to reduce optimistic estimates.

Functional testing should cover registration, login, profile management, prediction APIs, diet generation, chat responses, progress recording, dashboard retrieval, and database persistence. Integration testing should verify that the frontend can communicate correctly with the backend and that user-specific records are not mixed between accounts.

The Generative AI component can be evaluated for structural validity, relevance to the user's constraints, consistency, completeness, and safety. Generated meals can be reviewed against trusted nutritional references when available. Agentic functionality can be evaluated by testing whether changes in stored progress result in appropriate changes to recommendation context.

Usability evaluation can measure task completion, response clarity, perceived usefulness, and user satisfaction. A future study may involve students or general users performing representative tasks such as creating a profile, obtaining a diet plan, checking a prediction, recording progress, and asking the AI assistant a nutrition question.

Area	Measures
Disease models	Accuracy, precision, recall, F1-score, confusion matrix and ROC-AUC where appropriate.
API	Endpoint success rate, validation, error handling and response time.
Nutrition generation	Constraint satisfaction, structured-output validity and human/expert review.
Tracking	Record creation, retrieval, user association and chart consistency.
PDF reporting	Generation success, readability and completeness.
Usability	Task completion, navigation effort, clarity and satisfaction.
Agentic adaptation	Correct trend detection and appropriate recommendation-context updates.

Fig. 5. Recommended evaluation plan for the completed implementation.

XIII. RESULTS AND DISCUSSION

The integrated design of NutriGuardian AI creates an end-to-end workflow from user registration to personalized nutrition guidance and progress monitoring. The primary architectural benefit is shared context: the same user profile can support prediction, meal generation, conversation, and tracking rather than requiring separate data entry in each feature.

The disease-prediction module adds a preventive-health dimension to nutrition planning. Instead of generating a meal plan solely from weight or calorie goals, the system can incorporate predicted risk information as an additional context variable. This does not convert the system into a clinical diagnostic tool; rather, it provides a software-level risk signal that can be considered when generating supportive nutrition guidance.

The Generative AI component increases flexibility compared with a fixed meal database. It can explain recommendations and produce variations according to preferences. However, this flexibility introduces the need for validation because language models can generate nutritionally inappropriate or inconsistent information. The design therefore emphasizes structured prompting, output constraints, and future integration with trusted food-nutrient databases.

The agentic layer provides the mechanism for adaptation. A one-time recommendation is based mainly on the current state, whereas an agentic workflow can incorporate historical progress. For example, if the user's weight trend or recorded calorie intake changes over time, the system can use that change to update recommendation context. The usefulness of this feature depends on sufficient and accurate progress data.

The sustainability component broadens the scope of the platform. A nutrition system that considers health, affordability, local availability, and sustainability can support more practical recommendations. However, sustainability scoring must be transparent and based on defined data sources to avoid giving users misleading environmental claims.

Overall, the system demonstrates how ML, Generative AI, and agentic concepts can be combined within a modular web architecture. The final academic evaluation should report measured

model metrics, API test results, generation-quality observations, and user-study findings rather than illustrative values.

XIV. SCREENSHOTS



F

Fig. 6. Login and Register page

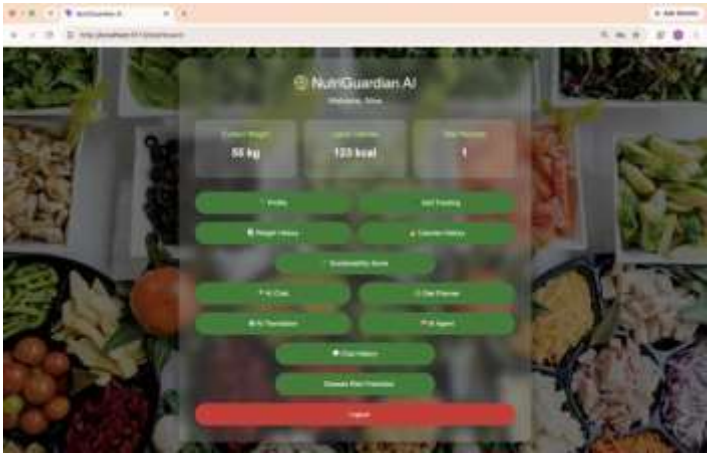


Fig. 7. Dashboard page

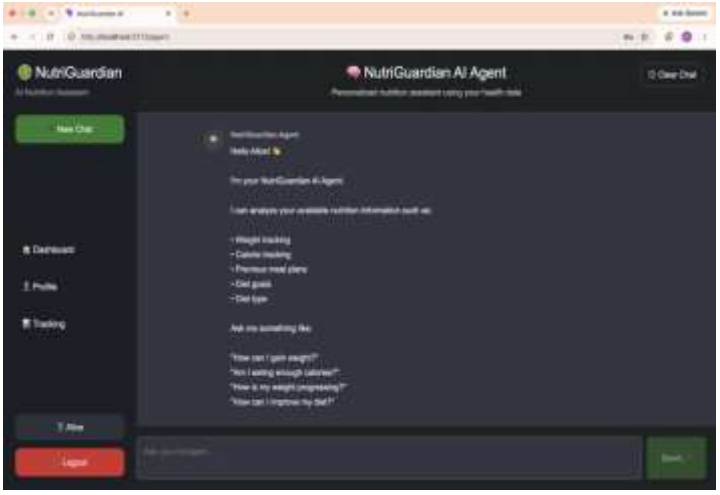


Fig. 8. NutriGuardian AI Agent page

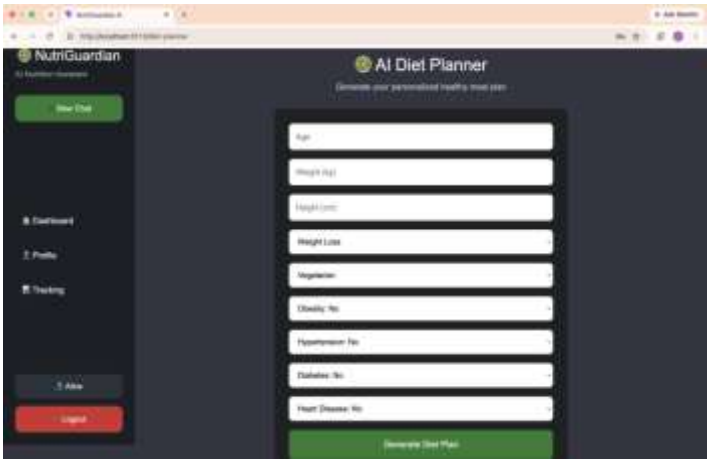


Fig. 9. AI Diet Planner page

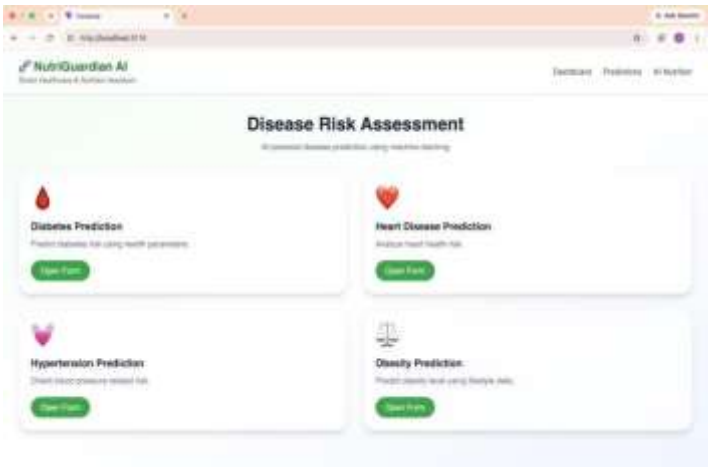


Fig. 10. Disease Risk Prediction page

XV. ADVANTAGES

The proposed system provides several functional advantages: personalized recommendations instead of a fixed diet chart; integration of disease-risk prediction with nutrition planning; conversational explanations through Generative AI; persistent user and progress history; adaptive feedback through agentic monitoring; dashboard-based visualization; region-friendly and affordability-aware meal planning; and sustainability-oriented nutrition support. The modular architecture also allows future models and services to be added without redesigning the entire application.

XVI. LIMITATIONS

The accuracy of disease-risk prediction depends on the quality, size, representativeness, and feature coverage of the training datasets. A model trained on a limited population may not generalize to every user. Generative AI may produce plausible but incorrect nutritional statements. Sustainability indicators depend on the quality and availability of environmental-food data. Agentic adaptation also depends on consistent progress records; sparse or incorrect user inputs can reduce its effectiveness.

The current system is a software prototype and should not be treated as clinically validated medical decision support. Clinical validation, expert review, larger and more diverse datasets, privacy assessment, bias evaluation, and prospective user studies are required before making clinical claims. These limitations should be explicitly considered when interpreting the system's outputs.

XVII. FUTURE WORK

Future work can include larger and more diverse datasets, comparison of multiple disease-prediction algorithms, calibration of risk outputs, trusted nutrition-database integration, automated nutrient validation, stronger prompt engineering, structured-output verification, and improved agentic policies. Wearable-device integration could provide activity and physiological information for more continuous personalization.

Additional extensions include food-image recognition, multilingual nutrition guidance, regional cuisine databases, grocery-cost optimization, personalized shopping lists, stronger sustainability metrics, explainable prediction dashboards, and mobile deployment. Human-in-the-loop review can also be introduced so that high-risk or ambiguous recommendations are flagged for professional attention.

A further research direction is to evaluate whether adaptive recommendations genuinely improve adherence and nutritional outcomes over time. Such evaluation should compare baseline and adaptive workflows using controlled user studies and clearly defined outcome measures.

XVIII. CONCLUSION

NutriGuardian AI presents an integrated approach to personalized nutrition and preventive healthcare by combining machine-learning-based health-risk prediction, Generative AI, agentic monitoring, progress tracking, and sustainability-oriented guidance. The system is designed as a modular web application in which user information is maintained as shared context across multiple services. The inclusion of diabetes, heart disease, hypertension, and obesity risk prediction extends the system beyond a conventional meal planner, while the Generative AI and agentic layers provide conversational and adaptive capabilities.

The proposed architecture demonstrates how prediction, generation, tracking, and feedback can be connected within one platform. Although the system is not a replacement for medical professionals and requires further validation, it provides a foundation for accessible and personalized nutrition support. Future research can strengthen model validation, nutritional verification, agentic decision policies, sustainability measurement, privacy, and user evaluation to improve reliability and practical usefulness.

XIX. REFERENCES

[1] D. Tanwar, T. Nobari, P. Chaparro, and A. Panangadan, "Quantitative Evaluation of AI-generated Recipes for Health Recommender Systems," Proc. 2025 IEEE Int. Conf. Information Reuse and Integration and Data Science (IRI), 2025, DOI: 10.1109/IRI66576.2025.00074.

[2] S. Shobana, S. Subitsha, R. Priscilla, and J. Gold Beulah Patturose, "Smart Diet Planner: An AI-Powered Web Application for Personalized Nutrition Management," Proc. 2025 Int. Conf. Future Technologies (ICFT), 2025, DOI: 10.1109/ICFT66708.2025.11336451.

[3] K. Samy, A. Osheba, H. Shams El-Din, R. M. Sallam, and M. Abd Elaziz, "Personalized Nutrition Recommendations Using Machine Learning," Proc. 2024 Int. Conf. Smart-Digital-Green Technologies and Artificial Intelligence Sciences (CSDGAIS), 2024, DOI: 10.1109/CSDGAIS64098.2024.11064829.

[4] M. Ge, F. Ricci, and D. Massimo, "Health-aware food recommender system," Proc. 9th ACM Conf. Recommender Systems, 2015, pp. 333–334.

[5] P. Chavan, B. Thoms, and J. Isaacs, "A recommender system for healthy food choices: building a hybrid model for recipe recommendations using big data sets," Proc. 54th Hawaii Int. Conf. System Sciences, 2021.

[6] Y. Tian, C. Zhang, Z. Guo, C. Huang, R. Metoyer, and N. V. Chawla, "RecipeRec: A heterogeneous graph learning model for recipe recommendation," arXiv:2205.14005, 2022.

[7] R. Yera, A. A. Alzahrani, and L. Martínez, "Exploring post-hoc agnostic models for explainable cooking recipe recommendations," Knowledge-Based Systems, vol. 251, p. 109216, 2022.

[8] Y. Chen, A. Subburathinam, C.-H. Chen, and M. J. Zaki, "Personalized food recommendation as constrained question answering over a large-scale food knowledge graph," Proc. 14th ACM Int. Conf. Web Search and Data Mining, 2021, pp. 544–552.

[9] M. Chen, X. Jia, E. Gorbonos, C. T. Hoang, X. Yu, and Y. Liu, "Eating healthier: Exploring nutrition information for healthier recipe recommendation," Information Processing & Management, vol. 57, no. 6, p. 102051, 2020.

[10] W. Wang, L.-Y. Duan, H. Jiang, P. Jing, X. Song, and L. Nie, "Market2Dish: Health-aware food recommendation," ACM TOMM, vol. 17, no. 1, pp. 1–19, 2021.

[11] I. Papastratis, D. Konstantinidis, P. Daras, and K. Dimitropoulos, "AI nutrition recommendation using a deep generative model and ChatGPT," Scientific Reports, vol. 14, 2024, DOI: 10.1038/s41598-024-65438-x.

[12] K. Stefanidis et al., "PROTEIN AI Advisor: A Knowledge-Based Recommendation Framework using Expert-Validated Meals for Healthy Diets," Nutrients, vol. 14, no. 20, p. 4435, 2022.

13.I. Orue-Saiz, M. Kazarez, and A. Mendez-Zorrilla, "Systematic review of nutritional recommendation systems," Applied Sciences, vol. 11, no. 24, p. 12069, 2021.