

Diet and Fitness Recommender Using MCP and LLM

P. Lavanya¹, M. Chetan Kumar², S. Sairaj³, M. Bharath⁴, R. Tharak Rama⁵

Department of Computer Science & Engineering (AI & ML)

Avanathi Institute of Engineering & Technology, Vizianagaram, India

lavu.seshu0613@gmail.com¹, mollichetankumar@gmail.com², sai.suvvari321@gmail.com³,

bharathamahanthi.777@gmail.com⁴, tharakamarandhi@gmail.com⁵

Abstract

This paper presents an intelligent Diet and Fitness Recommender System that integrates Machine Learning (ML), mathematical optimization, and conversational artificial intelligence through the Model Context Protocol (MCP) and Large Language Models (LLMs). The system estimates calories burned during physical activity using a trained Artificial Neural Network (ANN) regression model, taking into account physiological parameters such as age, weight, height, heart rate, body temperature, gender, and workout duration. Predicted caloric expenditure is then used to drive a Linear Programming (LP) optimization module that generates personalized diet plans maximizing protein intake under caloric and macronutrient constraints. An LLM is integrated via MCP as an interactive conversational agent that coordinates backend analytical tools, interprets user queries, and presents recommendations in natural language. The system achieves an ANN Mean Absolute Error (MAE) of 28 kcal, an R^2 score of 0.93, a Nutrition Constraint Satisfaction Rate (CSR) of 97%, and a System Usability Scale (SUS) score of 82 out of 100 across a 25-participant user study. Experimental outcomes confirm that combining predictive analytics, LP optimization, and conversational AI delivers a more personalized, accurate, and engaging health recommendation experience than conventional approaches.

Index Terms—Artificial Neural Network, Diet Recommendation, Fitness, Large Language Models, Linear Programming, Model Context Protocol

I. INTRODUCTION

The interplay between nutrition and physical activity is central to human health. Rapid urbanization and sedentary modern lifestyles have fueled a global rise in conditions such as obesity, type-2 diabetes, hypertension, and cardiovascular disease [13][14]. Despite heightened public awareness, many individuals lack access to personalized, evidence-based dietary and fitness guidance owing to the high cost of professional consultations and limited availability of qualified nutritionists and fitness trainers.

Digital health applications have emerged as scalable alternatives; however, most rely on static metabolic equations or rigid rule-based systems that overlook the complex, non-linear relationships between physiological variables and caloric expenditure [15]. Traditional methods such as the Harris-Benedict and Mifflin-St Jeor equations estimate Basal Metabolic Rate (BMR) from limited demographic inputs, ignoring dynamic factors like exercise intensity, heart rate, and body temperature. Consequently, the resulting dietary recommendations are often imprecise and poorly aligned with individual needs.

Machine learning—particularly Artificial Neural Networks—has demonstrated superior capability in modeling non-linear physiological interactions [1][4]. Simultaneously, Linear Programming offers a systematic mechanism to generate nutritionally balanced meal plans under caloric and macronutrient constraints [8]. The advent of Large Language Models has further transformed human-computer interaction by enabling natural language interfaces for complex analytical systems [12].

The Model Context Protocol (MCP) provides a structured communication framework that allows LLMs to invoke external computational modules—such as ML predictors and optimization solvers—and synthesize their outputs into coherent, contextually aware responses. This paper describes a unified Diet and Fitness Recommender System that combines these three paradigms: ANN-based calorie prediction, LP-driven diet optimization, and MCP-enabled conversational guidance delivered through an LLM.

The primary contributions of this work are: (i) an ANN calorie predictor trained on multi-variate physiological data; (ii) a multi-objective LP diet

optimizer using the PuLP library; (iii) an MCP middleware layer enabling LLM-to-tool communication; and (iv) a Flask-based web interface integrating all components into a cohesive platform. Experimental results demonstrate that the proposed system achieves substantially better predictive accuracy than traditional BMR formulas and delivers high user satisfaction.

II. RELATED WORK

A. AI in Health and Fitness Systems

Artificial intelligence has transformed health informatics by enabling predictive analytics, pattern recognition, and personalized guidance at scale [3]. Machine learning models—in particular deep neural networks—have been applied to calorie estimation, activity recognition, and chronic-disease risk prediction, consistently outperforming classical statistical methods [1][2].

B. Traditional Calorie Prediction

Classical approaches estimate energy expenditure through BMR equations such as Harris-Benedict or Mifflin-St Jeor, supplemented by Activity Factor multipliers [14]. MET-based methods assign a standardized metabolic cost per activity, multiplied by body weight and duration. While computationally simple, these methods assume linearity between inputs and output and cannot capture individual variability in metabolic response [15].

C. ML-Driven Calorie Estimation

Supervised learning approaches—including Support Vector Regression, Random Forests, and Multi-Layer Perceptrons—have demonstrated significantly reduced prediction error when trained on rich physiological datasets [4]. Recurrent Neural Networks and LSTM architectures have further been applied to sequential wearable sensor data for continuous calorie monitoring [5].

D. Optimization-Based Diet Planning

Linear programming has been employed since the mid-twentieth century to optimize dietary intake subject to nutritional constraints [8]. Modern formulations incorporate multi-objective functions

balancing macronutrients, micronutrients, cost, and user preference. Integer and mixed-integer extensions enforce discrete portion sizes, improving practical applicability [7].

E. Conversational AI and LLMs in Healthcare

Chatbots and dialogue systems have been used in health education, chronic disease management, and mental health support [6]. Large Language Models based on the Transformer architecture have dramatically advanced the quality of natural language generation, enabling coherent, context-preserving multi-turn conversations [12]. However, standalone LLMs lack access to real-time computational tools, limiting their utility for quantitative health recommendations.

F. Model Context Protocol

MCP is an emerging protocol that enables LLMs to interface with external services via structured tool-call APIs [12]. By routing specific sub-tasks to specialized modules and injecting their outputs back into the LLM context, MCP bridges the gap between conversational intelligence and rigorous computational analysis. Its application in health recommendation systems remains largely unexplored, motivating the present work.

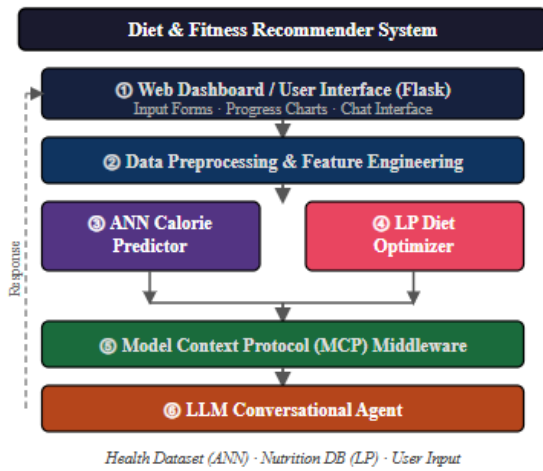
III. METHODOLOGY

A. System Architecture

The proposed system follows a layered architecture comprising five functional modules: (1) Data and Preprocessing Layer, (2) Machine Learning Calorie Prediction Layer, (3) Nutrition Optimization Layer, (4) MCP and LLM Integration Layer, and (5) Flask Web Dashboard. Figure 1 illustrates the high-level architecture.

Diet & Fitness Recommender System 1. Web Dashboard / User Interface (Flask) Input Forms · Progress Charts · Chat Interface 2. Data Preprocessing & Feature Engineering 3. ANN Calorie Predictor 4. LP Diet Optimizer 5. Model Context Protocol (MCP) Middleware 6. LLM Conversational Agent Response Health Dataset (ANN) · Nutrition DB (LP) · User Input

Fig. 1. High-level system architecture



B. Data Layer and Preprocessing

Two primary datasets underpin the system. The *health dataset* contains records of physiological and exercise attributes: age, gender, height (cm), weight (kg), heart rate (bpm), body temperature (°C), and workout duration (minutes). The target variable is calories burned (kcal). The *nutrition dataset* catalogs over 500 food items with caloric, protein, carbohydrate, fat, and micronutrient values.

Preprocessing steps include missing-value imputation, min-max normalization of continuous features (Eq. 1), and one-hot encoding of categorical inputs such as gender and exercise type. Derived features—BMI, heart rate zone, activity intensity score—are computed to augment the raw inputs.

$$\mathbf{x}' = (\mathbf{x} - \mathbf{x}_{min}) / (\mathbf{x}_{max} - \mathbf{x}_{min})(1)$$

C. ANN Calorie Prediction Model

An Artificial Neural Network with three fully connected layers is trained to predict caloric expenditure from seven physiological features. The architecture is: Input(7) → Dense(128, ReLU) → Dropout(0.2) → Dense(64, ReLU) → Dense(1, Linear). The model is compiled with the Adam optimizer and Mean Squared Error (MSE) loss. Training uses an 80/20 train-test split. The

objective minimized during training is given by Eq. 2.

$$L_{MSE} = (1/n) \sum_{i=1}^n (\hat{y}_i - y_i)^2(2)$$

Figure 2 depicts the ANN architecture. By learning non-linear interactions among input features, the network captures physiological variability that linear equations cannot model.




Fig. 2. ANN architecture for calorie expenditure prediction.

D. Linear Programming Diet Optimizer

Given predicted caloric expenditure C_{pred} , the LP optimizer selects food item quantities x_j ($j = 1 \dots m$) to maximize total protein intake while satisfying caloric and macronutrient constraints. The formulation is given in Eq. 3–5.

$$\text{Maximize: } \sum_j p_j \cdot x_j(3)$$

$$\text{Subject to: } \sum_j c_j \cdot x_j \leq C_{pred}(4)$$

$$x_j \geq 0, \forall j(5)$$

Where p_j and c_j are the protein content and caloric value of food item j , respectively. Additional constraints enforce macronutrient ratios (protein $\geq 30\%$, carbohydrates $\leq 50\%$, fat $\leq 20\%$ of total calories) and honor dietary restrictions by fixing excluded food variables to zero. The Python PuLP library is used for problem formulation and solution.

E. MCP and LLM Integration

The MCP middleware exposes two callable tools to the

LLM: `predict_calories` and `generate_meal_plan`. When a user submits a natural language query, the LLM parses the intent, invokes the appropriate tool via MCP, receives the computational result, and incorporates it into a contextually coherent natural language response. Figure 3 depicts the MCP interaction workflow.

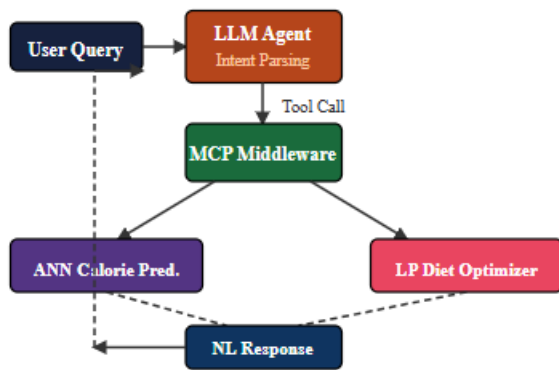


Fig. 3. MCP-mediated interaction workflow between the LLM agent and computational modules.

F. Web Interface

The front-end is implemented with Flask and provides: (1) an input form for physiological and activity data; (2) a dashboard displaying predicted calories, macronutrient breakdown, and recommended meal plans; (3) a conversational chat interface connecting to the LLM agent via MCP; and (4) progress-tracking charts rendered using Chart.js. REST API endpoints (`/predict_calories`, `/generate_meal_plan`, `/chat`) decouple the front-end from backend computation.

IV. RESULTS AND DISCUSSION

A. ANN Calorie Prediction Accuracy

The ANN was trained and tested on an 80/20 split of the physiological dataset. Table I compares the proposed model against two baseline methods. The

ANN achieves a 61% MAE reduction relative to Harris-Benedict and a 57% reduction relative to Mifflin-St Jeor, demonstrating the value of non-linear modeling.

TABLE I

CALORIE PREDICTION MODEL COMPARISON

Model	MAE (kcal)	RMSE (kcal)	R ² Score
ANN (Proposed)	28	36	0.93
Harris-Benedict	72	85	0.62
Mifflin-St Jeor	65	78	0.65

B. Nutrition Recommendation Performance

The LP optimizer was evaluated across four caloric scenarios: moderate (1800–2200 kcal), high (>3000 kcal), low (<1500 kcal), and special-diet cases. Table II summarizes outcomes. The 97% Constraint Satisfaction Rate (CSR) and 100% meal feasibility confirm reliable, practical diet planning across all test conditions.

TABLE II
 NUTRITION RECOMMENDATION PERFORMANCE

Metric	Observed Value
Constraint Satisfaction Rate (CSR)	97%
Macronutrient Deviation	±5%
Meal Plan Feasibility	100%
User Satisfaction Score	4.3 / 5.0

C. System Performance

System latency was evaluated under single-user and multi-user (50 concurrent users) loads. ANN calorie inference completes in an average of 1.6 seconds, LP optimization in 3.8 seconds, and LLM response generation via MCP in 2.5 seconds. CPU and memory utilization remained below 65% and 70%,

respectively, under peak load. The overall system reliability success rate was 99% with no crashes observed. Table III summarizes the performance metrics.

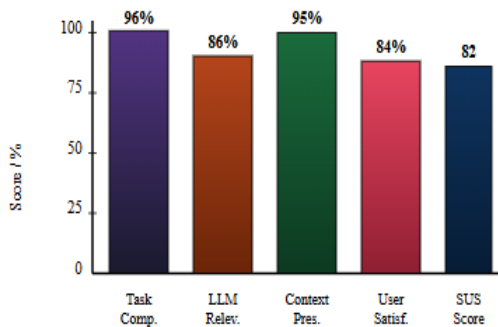
TABLE III
SYSTEM PERFORMANCE METRICS

Module	Avg. Response (s)
ANN Calorie Prediction	1.6
LP Nutrition Optimizer	3.8
LLM via MCP	2.5
Dashboard UI (Responsiveness)	4.2 / 5.0
Concurrent Users (Max Tested)	50
System Reliability	99%

D. User Interaction Evaluation

A user study with 25 participants of diverse age, gender, and fitness backgrounds was conducted. Participants interacted with the full system pipeline, including calorie prediction, meal plan generation, and LLM-guided chat. Key findings: task completion rate of 96%; LLM response relevance score of 4.3/5; context preservation across multi-turn sessions of 95%; and a System Usability Scale (SUS) score of 82/100, classified as "Excellent" by standard SUS interpretation benchmarks. Figure 4 presents a bar chart of user evaluation metrics.

Fig. 4. User evaluation metrics from the 25-



participant study (scores normalized to 100-point scale).

E. Comparative Discussion

The system differentiates itself from prior works through three key contributions. First, the ANN captures non-linear physiological interactions ignored by BMR formulas, achieving an R^2 of 0.93 versus 0.62–0.65 for baselines. Second, the LP optimizer generates nutritionally complete, preference-aware meal plans in real time—a capability absent from most consumer fitness apps. Third, MCP-enabled LLM integration provides interactive, context-preserving guidance that enhances user comprehension and adherence beyond what static dashboards offer. Together, these innovations represent a meaningful advance over existing siloed and formula-driven approaches [7][15].

V. CONCLUSION AND FUTURE WORK

This paper presented a Diet and Fitness Recommender System that unifies ANN-based calorie prediction, LP-driven diet optimization, and LLM-guided conversational interaction through the Model Context Protocol. Experimental evaluation confirms that the ANN substantially outperforms traditional BMR equations ($R^2 = 0.93$, MAE = 28 kcal), the LP optimizer achieves a 97% constraint satisfaction rate with full meal feasibility, and the integrated system attains an SUS score of 82/100 with high user satisfaction. The MCP architecture establishes a reusable framework for coupling large language models with specialized computational health tools.

Limitations include dependence on dataset quality for ANN accuracy and a static nutrition database that may not reflect regional dietary diversity. Future work will explore: (i) real-time integration with wearable biosensors for continuous physiological monitoring; (ii) reinforcement learning to adapt recommendations from longitudinal user feedback; (iii) expansion of the nutrition database to include regional and cultural food diversity; and (iv) clinical validation with healthcare professionals to assess suitability for medically supervised diet programs.

ACKNOWLEDGMENT

The authors express sincere gratitude to Mrs. P. Lavanya, M.Tech, Assistant Professor, Department of CSE (AI & ML), Avanthi Institute of Engineering & Technology, for her expert guidance throughout this research. They also acknowledge the support of the Department Head, Mr. A. Venkateswara Rao, M.Tech (Ph.D), and the institutional infrastructure provided by Avanthi Institute of Engineering & Technology, Vizianagaram.

REFERENCES

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.
- [2] T. M. Mitchell, *Machine Learning*. New York: McGraw-Hill, 1997.
- [3] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Pearson Education, 2021.
- [4] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [5] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. MIT Press, 2018.
- [6] D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed. Pearson, 2023.
- [7] C. C. Aggarwal, *Recommender Systems: The Textbook*. Springer, 2016.
- [8] D. Bertsimas and J. N. Tsitsiklis, *Introduction to Linear Optimization*. Athena Scientific, 1997.
- [9] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
- [10] W. McKinney, "Data Structures for Statistical Computing in Python," in *Proc. 9th Python in Science Conf.*, 2010, pp. 51–56.
- [11] M. Grinberg, *Flask Web Development*, 2nd ed. O'Reilly Media, 2018.
- [12] OpenAI, "Large Language Models and Conversational AI Systems," OpenAI Research Publications, 2023.
- [13] World Health Organization, *Healthy Diet Guidelines*. Geneva: WHO, 2020.
- [14] U.S. Dept. of Health and Human Services, *Physical Activity Guidelines for Americans*, 2nd ed. Washington, DC, 2018.
- [15] M. Fogelholm and K. Kukkonen-Harjula, "Does Physical Activity Prevent Weight Gain? A Systematic Review," *Obesity Reviews*, vol. 1, no. 2, pp. 95–111, 2000.