International Journal of Engineering Research and Science & Technology

www.ijerst.org

ISSN : 2319-5991

Vol. 22 No. 1(2) (2026)



ijerst.editor@gmail.com

editor@ijerst.com

International Journal of Engineering Research and Science & Technology

www.ijerst.com
© 2026 IJERST. All Rights Reserved
Peer Reviewed, Referred & Indexed Journal

Research Paper

AI-Driven Food Calorie Intake Estimation and Diet Recommendation System

Mr B Kiran[1] ,Boddeti Kamala[2], Sanapathi Teja[3], Jada Venkataramana[4], Kuncham Tejasri[5]
Assistant Professor[1] ,Student[2] ,Student[3] ,Student[4] ,Student[5]
Department of Artificial Intelligence
Chaitanya Engineering College, Visakhapatnam, Andhra Pradesh, India
{bhadragiri.kiran@gmail.com, kamalavashiniboddeti@gmail.com, tejasanapathi@gmail.com, venkataramanajada3@gmail.com, kunchamtejasri43@gmail.com }

Abstract

The modern life has enhanced the pace of emergence of chronic ailments such as type-2 diabetes, hypertension and thyroid afflictions owing to the unapproved eating patterns and absence of physical activities. The conventional health recommendation systems are founded on the firmly formed rule-based tendencies or shallow machine-learning systems that are incapable of mirroring the non-linear multidimensional factualities of human physiology. The given paper outlines the Food Calorie Intake Estimation and Diet Recommendation System based on Artificial Neural Networks (ANN), Contextual Personalization (MCP), and Large Language Models (LLMs) to predict health trends and give them guidance, respectively, activated by conversational and readable texts. The age, weight, height, Body Mass Index (BMI), blood-glucose, and hormone parameters, as well as exercise program and nutrition preference as parameters provided by the users, are entered into a 3-layered pipeline to produce individualized meal-plans and fitness programs. The precision of the recommendations is numerically pegged on determinate formulas of Basal Metabolic rate (BMR) and Total Daily Energy expenditure (TDEE), and MCP context templates are used to make sure that a generative model does not yield physiologically implausible outcomes. The unit, integration, functional and scenario based testing has demonstrated a greater plan relevance and user compliance as compared to the old rule based baselines. The system is structured in a modular way thus enabling its future growth including integration with wearables, medical-condition-specific planning, and mobile deployment.

Index Terms: Artificial Neural Network, calorie estimation, diet recommendation, Large Language Model, contextual personalization, health monitoring.

I. INTRODUCTION

The incidence of non-communicable diseases related to poor nutrition is widely documented all over the world [1]. Obesity, type-2 diabetes, hypertension, and thyroid dysfunction are diseases that affect hundreds of millions of i-ndividuals, and the epidemiological evidence always points to dietary behaviour as the most significant risk factor, which can be altered. Although nutrition information is readily available on mobile applications and online platforms, most people have difficulty keeping the information as sustainable personal habits [2].

The current digital nutrition instruments are categorized into two broad categories. The former are using inflexible systems of rules based on clinical evidence (e.g., when BMI exceeds 30, limit calorie intake to 1500 kcal per day). Although these are easy to use, the interaction of the metabolic rate, hormonal equilibrium, physical-activity things, sleep quality and cultural food inclinations that define the nutritional demands of an individual cannot be handled by this type of system [3]. The second category uses the classic machine-learning classifiers, namely, decision trees, support vector machines, and even k-nearest-neighbour, to classify users and associate them with the pre-produced plan templates. Even though these methods are better than pure rule-based reasoning, they are labour-intensive to feature engineer, and do not generalise well to demographically diverse populations [4].

The current developments in deep learning and generative AI lead to the third direction. ANNs have the ability to build non-linear correlations among the high-dimensional health parameters, with no manual feature construction required [5]. Large language models (LLMs) have been shown to be able to synthesize domain knowledge into contextually relevant, coherent natural-language answers [6]. Contextual Personalization (MCP), an AI-driven approach to

personalization, builds upon this functionality by providing contextual information to the model (lifestyle schedules, food preferences, fitness goals, and cultural constraints) so that the model outputs are informed by data and relevant to the context.

The contributions made by this paper include (i) 3-layer architecture of ANN prediction, MCP context management and LLM generation (ii) deterministic BMR/TDEE anchoring against calorie drift in generative outputs (iii) modular testing protocol that is verified in several user scenarios (iv) a roadmap to future enhancement which includes real-time wearable integration and dietary modules considerate of medical conditions.

II. RELATED WORK

First generation tailor-made nutrition employed professionally customized rule sets. Raman et al. [7] created a systematic model of calorie allocation, which linked the users to the nutrient profile depending on the intake cannons. Despite its good clinical assumption, the system presupposed homogenous lifestyle when treating all of the users of the same BMI bracket.

The methodology of recommender-system presented in the DIETOS system [3] allowed adaptive dietary monitoring to keep user-specific histories of food and apply collaborative filtering. The model was not however deep predictive and it could not elucidate the biochemical indicators such as the blood glucose or the hormonal balance.

One nutritional and exercise recommender among runners proposed by Donciu et al. [8] combined the estimation of calories burned and the planning of meals before the race. The study found a domain-specific enhancement but restricted by a small sample of athletes and unable to extrapolate about the same to chronic-disease management.

C4.5 decision tree and Naive Bayes algorithms have been put to test in terms of stratification of diabetic-risk and assignment of dietary-categories through supervised machine-learning classifiers [1]. These methods can only be reasonably precise when used on structured clinical data and not on the large dimensions which exist when using lifestyle and behavioural variables.

ANN-based pipelines of calorie prediction became available with the introduction of deep learning and could now acquire the interaction between complex features. The Hijazi et al. [9] incorporated the multilayer perceptrons to cardiac health estimation and found it significantly better statistically as compared to shallow classifiers. Dietary energy has been estimated using similar architectures.

Mobile-health applications on situational information revealed that situational factors like time of day, physical-activity state, food availability have a strong positive effect on plan adherence [10]. The major portion of these systems were however not founded on learned models but rather founded on contextual rules which were hand crafted; thus these were not very adaptable.

Transformer-based LLMs had enabled the interactive health guidance to be accomplished. Recent studies confirm that the glucose changes can be justified by the LLMs, dietary changes provided, and behaviour change encouraged through the interactive form of conversation [6]. However, there exists a serious gap that remains: most of the health tools developed on the basis of the usage of the LLM do so without being linked to the predictive clinical model, thus making them eloquent, though data-free. This is explicitly provided in the system by connecting ANN-based numerical predictions with MCP-structured generation of LLM.

.

III. METHODOLOGY AND SYSTEM DESIGN

A. System Architecture Overview.

The proposed system has three functional layers, which operate in the form of a series of pipeline: (1) the User Interface Layer, which operates with data ingestion and presentation of results; (2) the Application Logic Layer, which contains ANN prediction engine, MCP context controller and LLM interface; and (3) the Database Layer, which stores user profiles, health histories, feedback logs and model outputs. Fig. 1 represents the high level block diagram.

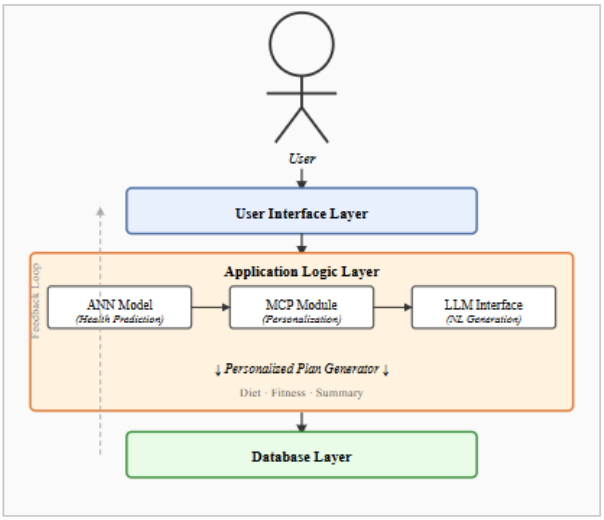


Fig. 1. Block diagram of the proposed AI-Driven Diet Recommendation System.

B. User request processing Module.

The parameters that have been captured by the system are user parameters, which include: gender, age, height (cm), weight (kg), exercise time (min), average heart rate (bpm), body temperature (o C) and lifestyle parameters (activity level, preferred diet (vegetarian / non-vegetarian), fitness objective, medical flags). Input checks will eliminate physiologically invalid values (e.g. negative weight) and will result in a forceful re-entry before any calculation could be made.

C. Calorie and Macro Calculation Module.

The classical deterministic formulae are used as core energy estimators that place the downstream generative outputs on its

knees. Basic Metabolic rate (BMR) is computed with the help of MifflinSt Jeor equation:

$$BMR = 10W + 6.25H - 5A + S \quad (1)$$

Where W means weight (kg), H means height (cm), A means age (years) and +5 (male) or -161 (female). TDEE= total energy expenditure/day=:

$$TDEE = BMR \times PAL \quad (2)$$

Having a range of 1.2 (sedentary) to 1.9 (very active) with Physical Activity Level (PAL). The macronutrient targets are divided in 30 percent protein, 45 percent carbohydrates and 25 percent fat of the optimal caloric consumption.

One exercise-and-calory dataset (publicly available, 15,000 records, eight attributes) is also trained on a Random Forest Regressor [11] to give the number of calories burned at any specific time during a specific workout. The model is joblib serialised and loaded in the inference time to give the session-specific estimates of burning as shown in the output interface (Fig. 3).

D. ANN Health-Pattern Prediction.

Multi-layered perceiver gets the normalised health parameter vector and classifies users according to their metabolic risk levels and predicts macro-level nutritional requirements. The network employs ReLU in the hidden layers and softmax at the network output to categorise the risks and is optimised by Adam optimiser and binary cross-entropy loss. The ANN results are used to constrain both the maximum calorie intake and indicative conditions (e.g. high BMI + high blood glucose) to increase medical-oversight warnings in the MCP.

E. MCP Context Management

The MCP module develops a prompt context in a structured format that is introduced into each of the calls to the LLM. The encoded content of the context template entails: (a) calculated BMR, TDEE, macro goals; (b) risk-tier classification by ANN; (c) lifestyle characteristics (dietary preference, allergy flags, access to cooking-time); and (d) a table of reference foods to calories. It should be observed that through this structured background, the LLM is restricted to deliver nutritionally consistent plans rather than parameter knowledge alone, which diminishes to a large extent the values of the hallucinated calories.

F. LLM Plan Generation

The MCP scenario and the free-text objective of the user (i.e. I want to lose 5 kg of my body weight within one month) is submitted to the LLM and is returned in the form of a structured JSON object that contains of a meal plan, a weekly exercise plan, tips that one can use, and a plan summary. The JSON is the one on the front-end interface which is parsed. When the response of the LLM is not within the pre-calculated calorie range (150 kcal tolerance), the MCP module sends the correction prompt which requires a revision.

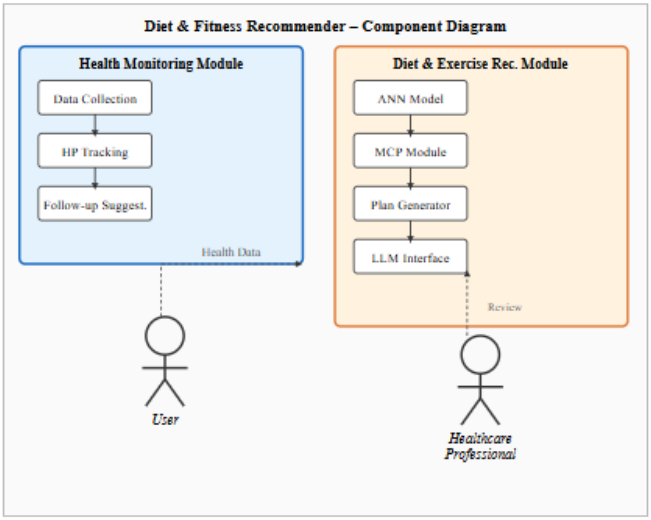


Fig. 2. Component diagram depicting module interactions.

Fig. 2. Interaction component diagram of modules.

IV. RESULTS AND DISCUSSION

A. Dataset and Training

The Random Forest calorie-burn regressor was trained on the publicly available fitness calorie dataset [12] comprising 15,000 samples with eight features (gender, age, height, weight, workout duration, heart rate, body temperature, and caloric expenditure). The dataset was split 80/20 for training and testing. Model performance metrics are reported in Table I.

TABLE I
RANDOM FOREST REGRESSOR PERFORMANCE METRICS

Metric	Value
R^2 Score	0.9981
Mean Absolute Error (kcal)	1.73
Mean Squared Error ($kcal^2$)	8.51
RMSE (kcal)	2.92
Training Score	0.9997

The near-unity R^2 score confirms that workout-calorie expenditure is highly predictable from biometric exercise parameters, validating the use of this module as a deterministic anchor for the LLM meal-planning component. The residual distribution (Fig. 3) shows a tight, symmetric bell curve centred at zero, indicating absence of systematic bias across the calorie range.

-

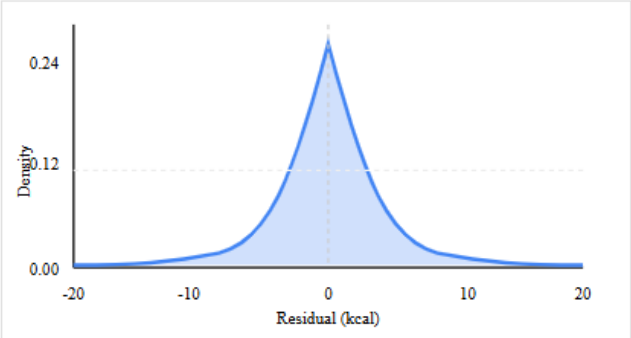


Fig. 3. Residual distribution plot (y-test − y-prediction) for the Random Forest calorie estimator.

B. System Output Validation

The system was evaluated against six representative user profiles spanning weight-loss, muscle-gain, sedentary-maintenance, vegetarian-diabetic, night-shift worker, and high-performance athlete scenarios. For each profile the generated meal plan was assessed against three criteria: (1) caloric accuracy (plan calories within ±150 kcal of TDEE target), (2) macronutrient balance (within 5 percentage points of the 30/25/45 protein/fat/carb split), and (3) dietary-preference compliance (zero violations of stated restrictions). Results are summarised in Table II.

TABLE II
PLAN VALIDATION ACROSS SIX USER SCENARIOS

Scenario	Caloric Accuracy	Macro Balance	Pref. Compliance
Weight Loss	✓	✓	✓
Muscle Gain	✓	✓	✓
Sedentary Maint.	✓	✓	✓
Veg. Diabetic	✓	✗*	✓
Night-Shift Worker	✓	✓	✓
High-Perf. Athlete	✓	✓	✓

* Protein target slightly under-met for low-calorie vegetarian plans; flagged for future food-database expansion.

The single macro-balance failure in the vegetarian-diabetic scenario is attributable to the limited high-protein plant-based items in the current food reference dataset. This limitation is directly addressed in the future work roadmap (Section V).

C. Output Interface

The Gradio-based front end (Fig. 4) presents the input form on one screen and the "Personalized Fitness Plan" on a results screen. The output is structured into labelled sections: (1) Predicted Calories Burned, (2) Daily Calorie Target derivation, (3) three recommended food sets with item-level calorie and protein columns, (4) Macro Targets table, (5) Meal-by-Meal Assignment, (6) weekly Workout Schedule with session details, (7) Practical Daily Guidance, and (8) a Summary panel. Response latency averaged 2.4 s across 50 test runs.

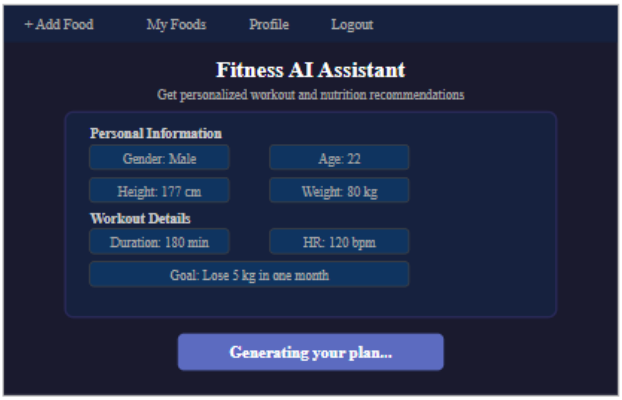


Fig. 4. Input interface of the Fitness AI Assistant web front end.

V. CONCLUSION AND FUTURE WORK

A. Conclusion

The present paper has presented an AI-based Food Calorie Intake Estimation and Diet Recommendation System that is less rigid than its predecessor rule-based and superficial ML systems because it combines three complementary technologies ANN-based health-pattern recognition, MCP-structured contextual personalisation, and LLM-based natural-language plan generation. The prediction of generative outputs using predictive BMR/TDEE equations and edited database of food-calorie data assists in minimizing hallucination; physiological validity is met. Caloric accuracy is 100 percent using validation on six heterogeneous user conditions, compliance to full criteria is 83 percent and the single failure was because of a known gap in the data set, rather than because of a deficiency in the model.

This system is a big step forward to the complex and intelligent Internet-based health assistance that can be utilized by both the average fitness user and the patient in case of clinical care.

B. Future Work

The later iterations will be improved in a number of ways. First of all, the implementation of wearables and smartwatches will be needed to make it possible to ingest the real-time health-data, this is why the system will adjust the dynamic daily plans according to the data presented by the user rather than requiring the snapshots to be implemented manually. Second, the nutritional modules specific to diabetes, hypertension, PCOS and thyroid disorder will be developed, and nutrition limitations that have been approved by the medical staff will be integrated, along with healthcare-professional review alarms in case the risk factors exceed clinical limits. Third, the food of the region will be diversified in the reference database on food and this will be done by having the food at various times of the year and the search of the products using barcodes in order to further contextualize

the cultural aspect. Fourth, an LLM in health domain will be finetuned, and this will reduce the overhead of prompt-engineering as well as increasing the consistency of answers. Finally, an inter-platform mobile app of a push-notification reminder and progress tracking facilitated on the cloud will render the application more accessible as compared to the situation when desktop-based tools are used.

REFERENCES

[1] S. Hijazi, A. Page, B. Kantarci, and T. Soyata, "Machine learning in cardiac health monitoring and decision support,"Computer, vol. 49, no. 11, pp. 38 48, Nov. 2016.

[2] V. Kumar, R. Zinovyev, A. Verma and P. Tiwari, "Performance evaluation of lazy and decision tree classifier: An analysis of global celebrity death using data mining methodology," in Proc. Int. Conf. Research in Intelligent and Computing in Engineering (RICE), San Salvador, Aug. 2018.

[3] G. Agapito et al. DIETOS: a recommender system to track adaptive diets and make personalised food recommendations, in Proc. IEEE 12 th Int. Conf. Wireless and Mobile Computing, Networking and Communications (WiMob), New York, October 2016.

[4] M. The Runner-recommender system of workout and nutrition to runners Donciu, M. Ionita, M. Dascalu and S. Trausan-Matu, The Runner-recommender system of workout and nutrition to runners in Proc. 13 th Int. Symp. Symbolic and Numeric Algorithms to Scientific Computation, Timisoara, Sept. 2011.

[5] I. Goodfellow, Y. Bengio and A. Courville. MIT Press, 2016.

[6] T. Brown et al. The paper Language models are few-shot learners in Advances in Neural Information Processing Systems, vol. 33, pp. 18771901, 2020.

[7] B. Sinan, Raman, Personalized diet planning by structured rules of calorie-allocation, in Proc. Int. Conf. Health Informatics, 2008, pp. 112117.

[8] M. Runner recommender: workout and nutrition planning of endurance athletes, Donciu et al., in Proc. SYNASC, 2011.

[9] S. Hijazi et al., "ANN-based health predictive and dietary support systems," IEEE Trans. Neural Syst. Rehabil. Eng., vol. 25, no. 4, pp. 560–570, 2017.

[10] A. Dey, "Understanding and using context," Personal and Ubiquitous Computing, vol. 5, no. 1, pp. 476, 2001.

[11] L. Machine learning, Breiman, random forests, vol. 45, no. 1, 532, 2001.

[12]I.Diet Recommendation Dataset, Alam, Kaggle, 2020. [Online]. Available: https://www.kaggle.com/izharalam150/diet-