Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

NutriSafe: An AI-Assisted Multi-Condition Food Safety Platform for Personalized Dietary Decision Support

Sarthak Chauhan
Department of Information Technology

Pranay bale
Department of Information Technology

Devendra Ghate
Department of Information Technology

Yash Jangle
Department of Information Technology

Abstract **-** Individuals with long-term dietary restrictions frequently need to determine whether a single food item complies with multiple medical regulations. The majority of nutrition apps display nutrient and calorie counts. The user is left to determine the medical interpretation. This gap is filled by NutriSafe. The food name or picture is entered into the system. It uses an AI vision service to identify the food. Standard nutrition data is pulled. It compares the information to a dietary and medical profile that has been stored. Diabetes, celiac disease, hypertension, chronic kidney disease, and food allergies are all covered by the prototype. One of three outcomes can be obtained from a fixed multi-condition safety engine: Safe, Moderation, or Not Recommended. It documents the components or nutrients that contribute to each outcome. Once the user verifies the extracted data, NutriSafe also allows prescription-based profile setup. It offers workable regional and Indian options. The fixed safety rules are kept apart from probabilistic AI tasks like document extraction and food identification. The ultimate decision remains dependent on well-defined guidelines and organized nutrition data. A functional prototype demonstrates the entire process, including local options, nutrient-level explanation, profile management, and food check. Quantitative food-recognition tests, broader nutrition coverage, and clinical review prior to any use outside research decision support are necessary for future work.

Keywords. *Food safety, computer vision, multi-condition reasoning, explainable AI, prescription OCR, personalized nutrition, chronic illness, dietary recommendations, Indian food, and React Native.*

I. INTRODUCTION

Choosing food becomes challenging when a single item needs to satisfy multiple medical requirements. A diabetic needs to monitor their carbohydrate intake. A person with chronic kidney disease or hypertension needs to restrict their intake of sodium and related nutrients. A person with celiac disease needs to stay away from gluten. Certain ingredients must be avoided by someone who has a food allergy. Nutrition therapy must be tailored to the individual's needs and preferences, according to the 2026 American Diabetes Association Standards of Care. Dietary modifications for CKD must be commensurate with the severity of the disease and other conditions, according to the 2024 KDIGO guideline [10], [9]. When choosing food in their daily lives, users outside of a clinic encounter an obvious information problem.

Meal tracking is made easier by generic nutrition apps. Automated diet assessment is supported by food images, according to computer vision systems. Convolutional neural networks were evaluated by DeepFood using Food-101 and UEC-256 [2]. NutriNet achieved 86.72 percent classification accuracy after training a deep model on over 225,000 photos in 520 food and drink categories [3]. Mobile multi-food work estimated calories and nutrients and identified multiple items in a single meal photo [4]. These studies demonstrate the effectiveness of food recognition. The question of whether a particular food fits a particular medical profile is not addressed by them.

Dietary advisors take a different tack when it comes to personalization. DIETOS uses medical questionnaires to create health profiles and suggests foods for chronic illnesses and regional cuisine [5]. Yum-me integrates learned food preferences, nutrition objectives, and dietary restrictions into a single meal pipeline [6]. The value of user profiles is displayed by these systems. NutriSafe concentrates on a single, integrated route from food identification to safety inspection. Before returning a recommendation, it simultaneously assesses multiple medical limits.

Indian cuisine creates a useful demand. Regional sampling of the Indian food supply is reflected in the Indian Food Composition Tables 2017, which provide nutrition information for 528 important foods [7]. For Indian foods, NutriSafe uses USDA FoodData Central as a secondary standard source and IFCT 2017 as its primary source [8].

This paper presents NutriSafe as a full decision-support system. The main contributions follow.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

- An integrated path for food identification and nutrition retrieval using text and images.
- a fixed safety engine that simultaneously verifies several dietary and medical restrictions.
- Explanations of nutrients and ingredients for every safety outcome.
- Setting up a prescription-based medical profile requires user approval.
- Indian and local food recommendations within the decision-making process..

NutriSafe is a prototype for research. A doctor or registered dietitian cannot be replaced by it. Dietary restrictions are based on the stage of the illness, prescription drugs, test results, portion sizes, and personal treatment regimens.

II. RELATED WORK

A. Food Image Recognition

Deep convolutional and transformer models replaced manually created visual features in food recognition. Food-101, which has 101,000 photos and 101 food categories, was published by Bossard et al. [1]. DeepFood tested on Food-101 and UEC-256 after applying CNN recognition to smartphone food photos [2]. Recognition and detection were handled by NutriNet, which reported 86.72 percent recognition accuracy and 94.47 percent detection accuracy [3]. Mobile multi-food photos and subsequent nutrition estimates were the focus of Pouladzadeh and Shirmohammadi [4].

Food logging is made easier by these systems. Recognition or diet assessment is their primary objective. A recognized food's clinical interpretation is still a different task.

B. Personalized Dietary Recommendation

User preferences, dietary objectives, or health information are used by personalized nutrition systems. Yum-me tested the strategy in user studies and used fine-grained food preferences to model dietary restrictions and nutrition goals [6]. DIETOS uses medical questionnaires to create a health profile and adds nutrition and disease information to a food catalog. CKD patients and healthy controls participated in the tests [5].

According to these studies, dietary recommendations need to be based on a long-term user profile. Immediate food-level safety inspections are made possible by NutriSafe. Before selecting an alternative, a user submits one dish and gets a condition-aware result.

C. Nutrition Databases and Clinical Guidance

The quality of the source data determines how well nutrition decisions are supported. An India-specific composition set comprising 151 food components for 528 essential foods is provided by IFCT 2017 [7]. Programmatic access and structured composition data are provided by USDA FoodData Central [8].

Clinical guidelines explain the failure of a single fixed threshold. KDIGO 2024 recommends that dietary modifications for CKD take into account protein, sodium, phosphorus, potassium, disease severity, and other factors [9]. Instead of focusing on a single macronutrient pattern, ADA 2026 emphasizes individual medical nutrition therapy [10]. Adult sodium intake should not exceed 2 g daily, according to WHO recommendations [11]. A strict gluten-free diet is mandated by celiac guidelines [12]. NutriSafe uses condition-specific guidelines instead of a single general health score because of these sources.

D. Research Gap

Four recurring limitations are identified in the literature: poor local context, weak multi-condition reasoning, generic nutrition guidance, and food recognition without medical personalization. Every limit is handled by NutriSafe in a single pipeline. The gaps are summarized in Table I.

RESEARCH GAP ADDRESSED BY NUTRISAFE

RESEARCH GAP	NUTRISAFE APPROACH
Generic calorie or nutrient information	Condition-aware safety result
Food recognition separated from health context	Food scan or text input linked to a stored medical profile
Single-condition advice	Multi-condition safety engine that checks all active limits together
Black-box advice	Nutrient and ingredient factors that explain the result
Limited local food context	Indian and regional food focus with local alternatives
Manual medical profile entry	Prescription extraction followed by user verification

Table. I. Research Gap Addressed

III. SYSTEM ARCHITECTURE

A layered design is used by NutriSafe. Fixed decision components and probabilistic AI components play different roles. Food and profile information are provided by the client. The food is identified and structured nutrition is retrieved by the processing layer. The profile is examined by the safety engine. The recommendation layer offers alternatives and provides an explanation of the outcome. The overall system architecture is depicted in Fig. 1.

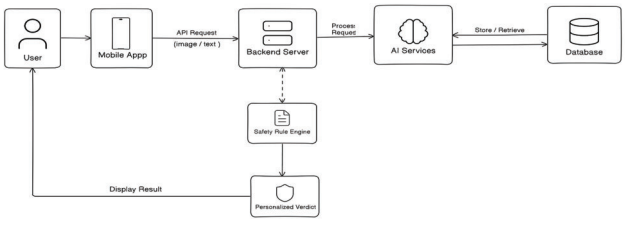


Fig. 1. Overall system architecture of NutriSafe.

A. Layer 1: Client Application

For cross-platform delivery, the client operates as a single React Native and Expo application. Food search, camera input, prescription upload, medical profile management, analysis history, and result display are all supported. Both mobile platforms are served by a single code base that maintains the camera workflow.

B. Layer 2: Input and Food Identification

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Two input paths are accepted by the system. A user uploads an image or enters the name of a food. An AI vision service returns the most likely food item or candidates from an image input. The system compares the name entered in text input to the food ontology. A canonical food identity connected to nutrition records is produced by the layer. The complete identification and safety pipeline is depicted in Fig. 2.

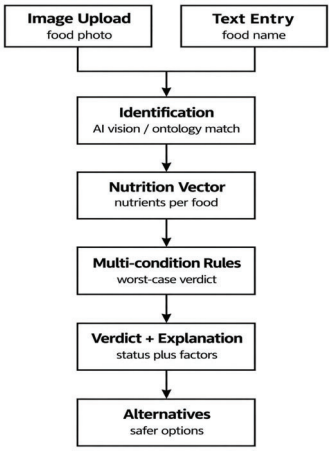


Fig. 2. End-to-end data pipeline from food input to safety verdict

Identification remains probabilistic. When it comes to similar dishes, mixed plates, preparation variations, or foods that are not in its training set, a recognition model faces uncertainty. The safety outcome is not determined by the identification stage.

C. Layer 3: Nutrition Retrieval

NutriSafe obtains nutritional data from structured databases following food identification. While USDA FoodData Central offers additional nutritional information for frequently consumed international foods, the Indian Food Composition Tables (IFCT 2017) give priority to Indian and regional foods.

Energy, carbs, protein, fats, sodium, potassium, phosphorus, fiber, sugar, gluten-related characteristics, and allergen information may all be included in the retrieved profile. For additional analysis, these values are normalized into a structured nutrition vector.

D. Layer 4: Medical Profile and Safety Engine

The user's current medical conditions, dietary restrictions, and allergies are stored in the medical profile. Diabetes, chronic kidney disease, hypertension, celiac disease, and food allergies are currently supported by NutriSafe. Subsequent food-checking sessions make use of this information.

The safety engine compares each active condition's rules to the nutrition vector. For instance, rules for diabetes take into account sugar and carbohydrates, rules for chronic kidney disease (CKD) look at potassium, phosphorus, sodium, and protein, and rules for hypertension mainly assess saturated fat and sodium.

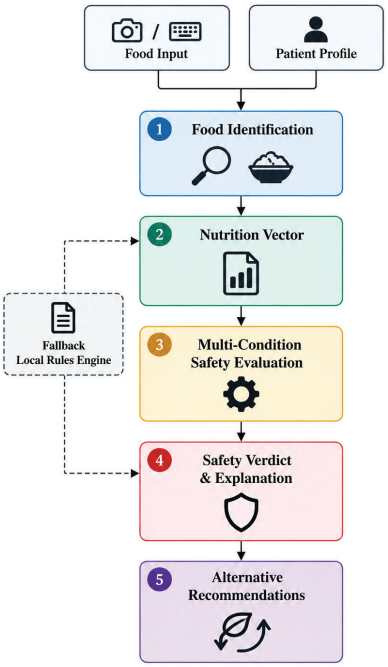


Fig. 3.Multi-condition safety engine decision flow.

The nutrition profile is compared to all active limits by the safety engine. An intermediate outcome is produced by each condition. All constraints are combined in the final product. If a food is prohibited by another active condition, it remains blocked.

E. Layer 5: Explanation and Alternatives

The result is one of three states: Safe, Moderation, or Not Recommended. In order to help users understand the rationale behind the recommendation, the app shows the ingredient or nutrient factors that contributed to the outcome. Each factor carries a level (Low, Moderate, High, Contains, or None) and an impact tag (positive, neutral, warning, or danger) with a short clinical detail, so the user sees not just the verdict but exactly which nutrients drive it.

When a food fails or needs to be moderated, the recommendation layer returns other foods that fit the same profile and favors sensible Indian and regional options. These substitutes are chosen based on the user's current conditions and pertinent dietary restrictions, which makes the suggestions more useful and customized.

Explanation quality is enforced at both generation and normalization time. The analysis request instructs the model to return a fixed schema a one-to-two sentence summary, a two-to-four sentence clinical justification, two to five nutrient factors, two to three practical swaps, and concrete portion guidance and a normalization pass then repairs the output.

F. Layer 6: Data, Authentication, and History

The prototype uses convex for data management and synchronization. It has a track record of long-term authentication and analysis. Optional caching reduces the

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

number of repeated data calls. Compared to standard app preferences, prescription and medical profile data require stricter privacy controls. Extracted prescription data never modifies a health profile without the consent of the user. Fig. 4 shows the workflow for the prescription-assisted profile.

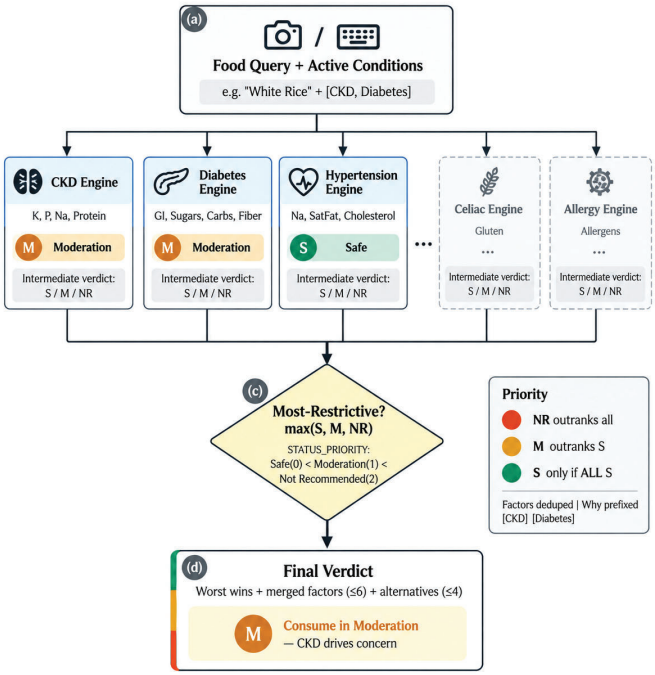


Fig. 4. Prescription-assisted medical profile initialization with mandatory user confirmation.

IV. METHODOLOGY

A. Knowledge Sources and Dataset Strategy

The approach combines three categories of knowledge: clinical guidance, nutrition databases, and food-recognition datasets. A general recognition benchmark is provided by Food-101. Local dish coverage is enhanced by regional or Indian datasets. The India-oriented nutrition reference is provided by IFCT 2017. Additional standard records are provided by USDA FoodData Central [1], [7], [8].

Clinical guidelines remain separate from the AI model. KDIGO 2024 specifies that dietary modifications must correspond with the severity of the disease and other conditions, and it provides CKD practice points on sodium, phosphorus, potassium, and protein [9]. Individual medical nutrition therapy for diabetes is emphasized by ADA 2026 [10]. Additional reference points are provided by WHO sodium guidelines and celiac guidelines [11], [12]. These sources are regarded as knowledge references in the prototype. It does not encode a single, universal medical threshold for each user.

B. Food Identification Pipeline

The food-identification pipeline goes from I to P (I to F). The input image is I. The processed image is P(I). F is the candidate set or anticipated food identity. A parallel path connects the query to the food ontology for text input.

The identified food is normalized by the system to a canonical identifier. The same food is referred to by several names or regional variations. The nutrition records are joined by the canonical identifier.

C. Nutrition Retrieval and Normalization

The nutrition service returns vector N(F) equal to energy, carbohydrate, protein, fat, sodium, potassium, phosphorus, allergen attributes, and gluten attributes for identified food F. As the body of knowledge increases, so does the field set.

Different serving definitions and units are used in nutrition databases. Units are normalized by the system prior to rule checks. Instead of appearing as zero, missing values are shown as explicit missing data. Instead of making an incorrect, precise medical decision when a necessary nutrient is missing, the system returns an uncertainty state or requests more information.

D. Multi-Condition Safety Model

Assume that the set of active conditions equals M, the user's active medical profile. A set of rules is defined by each condition. Each condition yields an intermediate result v in the set Safe, Moderation, Not Recommended for nutrition vector N. The outcome adheres to a rigid sequence: Safe is ranked lower than Moderation and Not Recommended. The overall outcome cannot be safe if one active condition returns Not Recommended. The most-restrictive constraint principle is put into practice in this way.

E. Prescription-Assisted Profile Initialization

The prescription image is D, and the prescription flow is D to OCR(D) to X to V(X) to M. Text is extracted by OCR. Candidate medical entities are stored in X. V stands for user confirmation. The medical profile is only filled out by verified entities.

A safety precaution is user confirmation. Handwriting, acronyms, dosage information, and irrelevant text can all be found in prescription documents. They may be misinterpreted by an extraction model. AI is used by NutriSafe to help with profile creation. Text that has been extracted is never regarded as final authority.

F. Explanation Generation

The engine logs the binding constraints for every outcome. Structured evidence provides the explanation. It doesn't ask a generative model to come up with a medical explanation. Condition, attribute, value, rule, and outcome are all contained in an explanation tuple. The nutrient or ingredient that triggered a condition-specific rule and the corresponding result state are indicated by the system.

G. Alternative Recommendation

The same safety engine is applied to alternative foods. The system retains candidates whose safety result is not Not Recommended if F is the original food and A is the candidate

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

set. Regional availability, food category, preference, and nutrition similarity are then used for ranking. The prototype prioritizes local and Indian options over hypothetical nutrient swaps.

V. IMPLEMENTATION AND PROTOTYPE

The prototype is a cross-platform TypeScript/JavaScript application using React Native with Expo, Convex for data synchronization, optional Redis caching, PostgreSQL/vector storage, and OAuth-style authentication, designed around separated client workflows, AI services, nutrition knowledge, fixed rules, and user data, with interfaces enabling food searches/scans, displaying verdicts and nutritional factors, storing medical profiles, and presenting key screens in Figures 5 to 8.

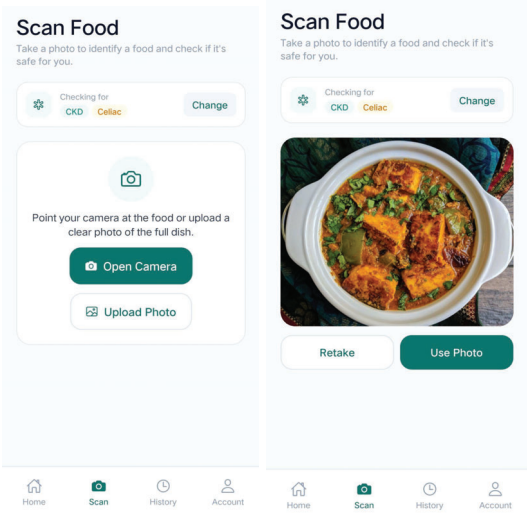


Fig. 5. Prototype food-check screen

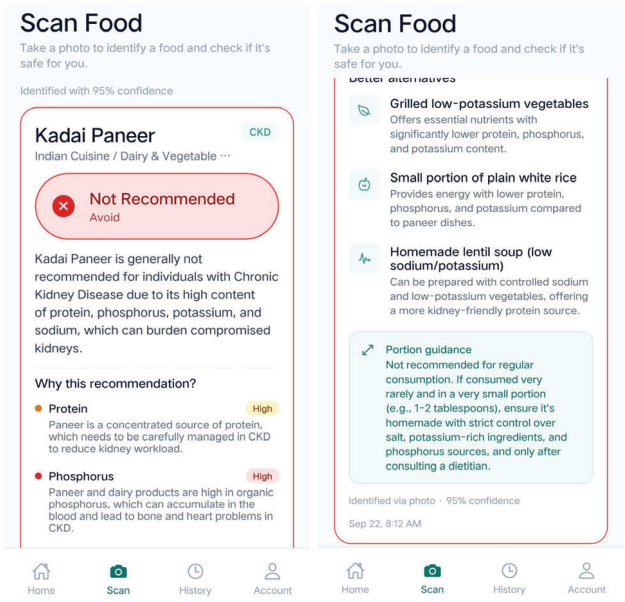


Fig. 6. Result screen showing tri-state verdict and nutrient-level explanation.

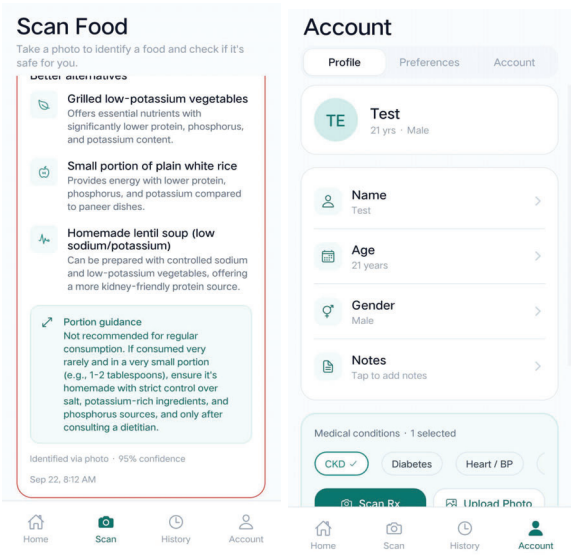


Fig. 7. Alternative recommendation list and Account screen

A. Software Requirements

The main categories of technology are listed below. Instead of being tested services, items that are marked as optional are implementation choices. The components are listed in Table II.

CORE SOFTWARE COMPONENTS

COMPONENT	ROLE IN NUTRISAFE
React Native plus Expo	Cross-platform mobile client
TypeScript and JavaScript	Application and service code
Convex	Profile, history, synchronization, and data operations
AI vision service	Image-based food identification
OCR and NLP service	Prescription text and medical-entity extraction
IFCT 2017	Primary India-oriented nutrition reference
USDA FoodData Central	Additional standard nutrition source
Deterministic rule engine	Multi-condition safety evaluation
Redis (optional)	Caching for repeated lookups
PostgreSQL or vector storage (optional)	Structured and semantic data extensions
OAuth-based authentication	User and profile access

Table. II. Placeholder for quantitative evaluation results (top-1/top-5 accuracy, latency, or confusion matrix).

VI. EXPERIMENTAL EVALUATION

A. Evaluation Scope

A functional prototype flow is displayed in the project materials. For the vision model, prescription OCR, or clinical safety classification, it does not provide a repeatable quantitative benchmark. No fake accuracy figures are reported in this paper. The assessment identifies the quantitative tests needed for the subsequent phase and concentrates on functional integration.

Food name and image input, food identification and nutrition retrieval, medical-profile matching, multi-condition rule evaluation, three-state result generation, nutrient-level explanation, alternative recommendation, prescription-assisted

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

profile setup, history and profile persistence are all displayed in the prototype.

B. Illustrative Safety Evaluation

Kadhai Panner is compared to a CKD profile in one demonstration. Along with other foods, the prototype shows potassium, sodium, and phosphorus as contributing factors and returns Safe to Consume. Instead of just a list of calories, the case illustrates the data flow from a food item to a condition-aware explanation.

C. Prototype-Level Functional Evaluation

The prototype's functional capabilities and the metrics intended for the subsequent evaluation phase are compiled in Table III.

EVALUATION CAPABILITY SUMMARY

CAPABILITY	NEXT-STAGE EVALUATION METRIC
Food text input	Input-to-food matching accuracy
Food image input	Top-1 and top-5 accuracy
Nutrition retrieval	Retrieval coverage and missing-field rate
Medical profile with conditions and allergies	Profile extraction accuracy
Multi-condition evaluation via fixed rule engine	Verdict agreement rate
Three-state decision (Safe, Moderation, Not Recommended)	Precision, recall, and F1 by class
Explanation with nutrient-level factors	Explanation correctness and user rating
Alternatives with safer options	Constraint satisfaction rate
Prescription extraction with user review	Entity precision and recall
End-to-end latency in prototype	Median and 95th-percentile latency

Table. III. Evaluation Capability Summary

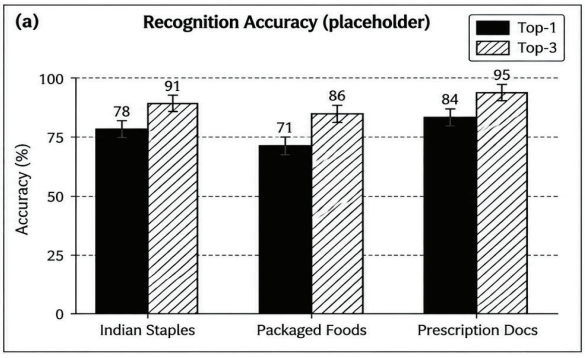


Fig. 8(a). Placeholder for recognition accuracy results (top-1/top-5 accuracy and confusion matrix).

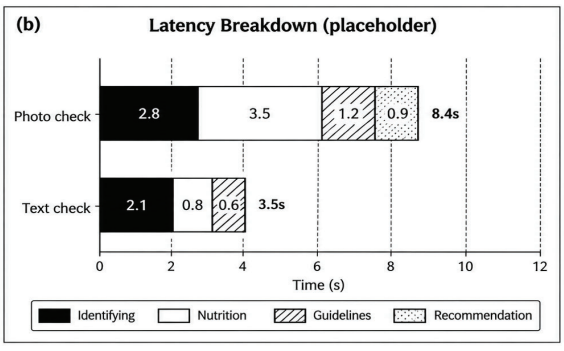


Fig. 8(b).Placeholder for system latency results (median and 95th-percentile)

D. Quantitative Evaluation Protocol

A comprehensive evaluation necessitates a held-out set of images of Indian food featuring both single-item and multi-item plates. Recognition metrics include top-1 and top-5 accuracy, precision, recall, and F1. A confusion matrix covers visually similar dishes.

The safety engine must be tested separately from the recognition model. A test suite combines medical profiles with conflicting conditions and food records. Safety agreement is the proportion of cases in which the system verdict concurs with the expected expert or rule-derived verdict. The precision, recall, and F1 for each of the three result states provide more details.

Prescription extraction requires entity-level precision, recall, and F1 for condition and allergy extraction. False-positive additions are significant because subsequent dietary decisions are influenced by an incorrect medical condition.

Food recognition, nutrition retrieval, rule evaluation, and the complete end-to-end request all require median and 95th-percentile latency for system performance. These figures distinguish between database and application overhead and AI inference time.

System performance requires median and 95th-percentile latency for food recognition, nutrition retrieval, rule evaluation, and the full end-to-end request. These numbers differentiate between AI inference time and database and application overhead.

Evaluation is repeatable rather than anecdotal because each of these layers corresponds to an automated gate in the project's continuous-integration pipeline. While end-to-end suites spin up a mock vision provider alongside the real server to assert full-request correctness for identification, nutrition analysis, and both multipart and base64 upload paths, including negative cases (missing token → 401, missing payload → 400, oversized image → 413), unit-level smoke tests test the AI layer against mocked providers. Each suite operates on every push and pull request, and the end-to-end harnesses log the timings of each request, providing an integrated latency log that can be combined over multiple runs to create the median and tail statistics mentioned above.

Robustness cannot be inferred from clean-lab accuracy; it must be measured explicitly. Thus, degraded-input ablations are part of the protocol: HEIC versus JPEG encodings, motion-blurred and low-light photos, partial dishes, and non-food photos, confirming that each degrades to a clear, actionable message (e.g., retake-as-JPEG guidance or an "unreadable document" state) instead of a confident misclassification. The deterministic local rules engine must replicate the same verdict schema in a second ablation that covers the offline fallback path with the AI provider disabled. In this mode, safety agreement is recalculated to measure the amount of accuracy retained without inference. Prior to the study, acceptance thresholds (such as top-1 recognition ≥ 0.85

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

with model confidence ≥ 0.72, safety agreement ≥ 0.95, prescription false-positive rate ≈ 0) should be established. Latency budgets should be broken down by stage in relation to the system's design envelope, which calls for a 45-second server-side vision timeout within a 90-second client budget.

E. Limitations of the Current Evaluation

Clinical validity, prospective patient outcomes, and agreement with registered dietitians on a large test set are not established by the existing material. A controlled held-out dataset is not used to report food identification performance. A prototype case study is the Kadhai Paneer outcome. It does not prove that all foods or patients are clinically correct.

Another restriction is portion size. Quantity, preparation, ingredients, and recipe all affect nutrition values. A food image can identify a dish, but it may not accurately reflect the amount consumed. Future tests must distinguish between portion estimation and recipe uncertainty and food identity accuracy.

Another restriction is portion size. Quantity, preparation, ingredients, and recipe all affect nutrition values. A food image can identify a dish, but it may not accurately reflect the amount consumed. Future tests must distinguish between portion estimation and recipe uncertainty and food identity accuracy.

VII. COMPARATIVE DISCUSSION

The primary way that NutriSafe is different from earlier systems is through integration. Food identification and diet data estimation are the main goals of food recognition systems like Food-101 methods, DeepFood, NutriNet, and mobile multi-food work [1], [2], [3], [4]. Yum-me and DIETOS are examples of personalized recommenders that concentrate on recommendations, health profiles, or preferences [5], [6]. NutriSafe integrates a fixed safety layer with these directions. A recognized food's suitability for a multi-condition profile is determined by the layer. The feature sets are compared in Table IV.

FEATURE-LEVEL COMPARISON OF RESEARCH DIRECTIONS

FEATURE	FOOD RECOGNITION SYSTEMS	MEAL RECOMMENDERS	DIETOS	NUTRISAFE
Food image recognition	Yes	Optional	No primary focus	Yes
Personal health profile	Limited	Yes	Yes	Yes
Nutrition retrieval	Yes	Yes	Yes	Yes
Multi-condition safety logic	Not primary	Limited	Condition-aware	Explicit
Three-state safety verdict	No	No	No	Yes
Nutrient-level explanation	Limited	Variable	Profile-based	Explicit
Indian and regional food focus	Dataset-dependent	Dataset-dependent	Regional	Indian and regional
Prescription-assisted profile	No	No	Questionnaire	Yes

Table. IV. Feature-Level Comparison of Research directions

It is a feature-based comparison. It does not assert that a particular line of inquiry is always superior. The objective is to identify the architectural gap that NutriSafe closes by combining food identification, structured nutrition, multi-condition limits, explanation, and local alternatives into a single user flow.

VIII. SAFETY, PRIVACY, AND RESPONSIBLE USE

Health-related data is processed by NutriSafe. Uncertainty and privacy are essential design criteria. User confirmation is required for medical conditions that are extracted from prescriptions. Profiles and analysis history are safeguarded by authentication and access control. Sensitive information is only retained for as long as the system needs it.

The safety engine distinguishes between information derived from an AI model and information from a reliable nutrition database. A dish may be incorrectly identified by a vision model. A prescription may be misread by an OCR model. Each recipe has a different nutritional value. The interface never displays a probabilistic model output as a clinical diagnosis and only displays uncertainty when necessary.

The prototype is intended for educational and dietary decision support applications. Physicians, nephrologists, endocrinologists, gastroenterologists, allergists, and registered dietitians are not replaced by it. Dietary restrictions for CKD are based on lab results and disease stage. Individual treatment objectives and medication regimens determine diabetes nutrition plans [9], [10].

IX. FUTURE SCOPE

The work is strengthened by a number of extensions. First, on larger Indian-food datasets that include mixed plates and regional dishes, the food-recognition component needs to be assessed and improved. Second, nutrition values must be mapped to the actual serving when estimating portion sizes. Third, recipe-level ingredient information and uncertainty ranges need to be added to the nutrition knowledge base.

More conditions and clinician-configurable limits can be covered by the multi-condition engine. Clinical guidance modifications can be audited independently of application code thanks to a versioned rule repository. Nutritionists could examine the evidence supporting a result and, if necessary, override or annotate a recommendation using a clinician-facing review interface.

Another approach is privacy-preserving education. Federated or on-device learning can enhance models without requiring the central gathering of unprocessed food or medical data. Under appropriate governance and validation, interoperability with health-data standards like FHIR is a later option.

A human-centered assessment is necessary. Future research must assess whether nutrient-level explanations enhance user comprehension, whether the options seem realistic and culturally appropriate, and whether the system reduces decision-making time without increasing risky confidence.

X. CONCLUSION

NutriSafe, a cross-platform AI-assisted system for condition-aware food safety and individualized dietary decision support, was introduced in this paper. The system combines text and image food identification, structured nutrition retrieval, a stored medical profile, prescription-assisted profile setup, fixed multi-condition safety rules, nutrient-level explanations, and recommendations for Indian and regional alternatives. The fundamental design

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

principle distinguishes fixed safety logic from probabilistic AI. AI helps with profile extraction and food recognition. Explicit guidelines and organized nutrition data provide the final dietary classification.

The complete end-to-end flow is depicted in the prototype, including food inspection, medical profile management, a CKD-focused kadhai paneer evaluation, nutrient-level reasoning, and alternative suggestions. Clinical validity and quantitative model performance are not established by the available data. Controlled food recognition benchmarks, rule-engine validation against expert-labeled cases, prescription extraction assessment, latency measurement, and user studies are all necessary components of the next research phase. Beyond general nutrition tracking, NutriSafe provides a foundation for explainable, localized, multi-condition dietary decision support with those assessments and clinical supervision.

REFERENCES

[1] L. Bossard, M. Guillaumin, and L. Van Gool, "Food-101 Mining discriminative components with random forests," in *Proc. European Conf. Computer Vision (ECCV)*, 2014, pp. 446-461, doi: 10.1007/978-3-319-10599-4_29.

[2] C. Liu, Y. Cao, Y. Luo, G. Chen, V. Vokkarane, and Y. Ma, "DeepFood: Deep learning-based food image recognition for computer-aided dietary assessment," *arXiv:1606.05675*, 2016.

[3] S. Mezgec and B. Korousic Seljak, "NutriNet: A deep learning food and drink image recognition system for dietary assessment," *Nutrients*, vol. 9, no. 7, p. 657, 2017, doi: 10.3390/nu9070657.

[4] P. Pouladzadeh and S. Shirmohammadi, "Mobile multi-food recognition using deep learning," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 13, no. 3s, Art. no. 36, 2017, doi: 10.1145/3063592.

[5] G. Agapito, M. Simeoni, B. Calabrese, I. Care, T. Lamprinoudi, P. H. Guzzi, A. Pujia, G. Fuiano, and M. Cannataro, "DIETOS: A dietary recommender system for chronic diseases monitoring and management," *Computer Methods and Programs in Biomedicine*, vol. 153, pp. 93-104, 2018, doi: 10.1016/j.cmpb.2017.10.014.

[6] L. Yang, C.-K. Hsieh, H. Yang, J. P. Pollak, N. Dell, S. Belongie, C. Cole, and D. Estrin, "Yum-me: A personalized nutrient-based meal recommender system," *ACM Trans. Inf. Syst.*, vol. 36, no. 1, Art. no. 7, 2017, doi: 10.1145/3072614.

[7] T. Longvah, R. Ananthan, K. Bhaskarachary, and K. Venkaiah, *Indian Food Composition Tables 2017*. Hyderabad, India: National Institute of Nutrition, Indian Council of Medical Research, 2017.

[8] U.S. Department of Agriculture, "FoodData Central," Agricultural Research Service. [Online]. Available: https://fdc.nal.usda.gov/. Accessed: Sep. 17, 2026.

[9] KDIGO, "KDIGO 2024 clinical practice guideline for the evaluation and management of chronic kidney disease," *Kidney International*, 2024.

[10] American Diabetes Association Professional Practice Committee, "5. Facilitating positive health behaviors and well-being to improve health outcomes: Standards of Care in Diabetes 2026," *Diabetes Care*, vol. 49, Suppl. 1, 2026, doi: 10.2337/dc26-S005.

[11] World Health Organization, "Sodium reduction," May 11, 2026. [Online]. https://www.who.int/news-room/fact-sheets/detail/sodium-reduction. Accessed: Sep. 17, 2026.

[12] J. M. McDermid, M. A. Almond, K. M. Roberts, E. M. Germer, M. G. Geller, T. A. Taylor, R. C. Sinley, and D. Handu, "Celiac disease: An Academy of Nutrition and Dietetics evidence-based nutrition practice guideline," *J. Acad. Nutr. Diet.*, vol. 123, no. 12, pp. 1793-1807.e4, 2023, doi: 10.1016/j.jand.2023.07.018.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)