# Label AI: Barcode Scanning based Mapping of Nutritional Values to Fitness Planning

Aditya Jha
Department of Artificial Intelligence & Data Science
Thakur College of Engineering and Technology
Mumbai, India

Tejas Gadekar
Department of Artificial Intelligence & Data Science
Thakur College of Engineering and Technology
Mumbai, India

Swati Joshi
Department of Artificial Intelligence & Data Science
Thakur College of Engineering and Technology
Mumbai, India

## ABSTRACT
Unhealthy dietary patterns and excessive intake of processed foods are major contributors to the global rise of obesity and chronic diseases, highlighting the need for accessible tools that enable consumers to make informed food choices at the point of purchase. Label-AI is a web-based system designed to address this challenge by scanning product barcodes by scanning a product's Universal Product Code (UPC) with a smartphone, LabelAI retrieves detailed nutrient data from an extensive food database (Open Food Facts) extracting nutrition information and generating a NutriScore-style health rating on a scale of 0-10.

The system's engine processes the nutritional information obtained from barcode scans and computes rating on a 0–10 scale based on key nutrients such as sugars, fat, saturated fat, salt, proteins, fiber, and energy per 100 g. Products with lower scores trigger alerts and suggestions for healthier alternatives within the same category. This paper presents the design and evaluation of Label-AI, including an overview of existing barcode-based nutrition applications, a two-tier architecture that combines browser-side scanning with cloud-based data retrieval, and a hybrid scoring mechanism that integrates machine learning with rule-based thresholds.

## Keywords
Barcode scanning; Nutrition tracking; Mobile health (mHealth); Personalized diet; Artificial intelligence; Machine Learning; Rule-based Scoring; Open Food Facts.

## 1. INTRODUCTION
Worldwide obesity and diet-related non-communicable diseases have reached alarming levels. As of 2023, roughly 38% of the global population was classified as overweight or obese, a figure projected to exceed half of humanity by 2035. Poor dietary behaviors—particularly high consumption of processed foods rich in fats, sugars, and salt—are major contributors to this epidemic.

LabelAI is designed to bridge this gap by combining the convenience of barcode scanning with AI-powered personalized nutrition guidance. The core idea is straightforward: whenever a user scans a packaged food, the app will log the item, evaluate it and provide a NutriScore on scale a of 0–10, and also recommend alternatives for the same.

By analyzing a product's nutritional profile, LabelAI can instantly indicate whether the item is generally healthy (a "green flag") or less advisable (a "red flag") through its NutriScore-style 0–10 rating. For instance, scanning a high-sodium instant soup may result in a low score and an alert such as "High sodium content – consider a healthier alternative." Similarly, a sweet candy bar with excessive sugar would be flagged with a warning and accompanied by suggestions for better options within the same category. While the current system focuses on population-level nutrition scoring and recommendations, future extensions will incorporate personalized thresholds, allergen detection, and diet-plan integration. On the positive side, the app can highlight nutritional details (e.g., "High Fructose Corn Syrup is linked to potential health concerns like increased risk of Type-2 Diabetes") to reinforce good and bad choices. This contextual feedback loop transforms passive logging into an active decision support system. By leveraging an extensive nutrition knowledge base, LabelAI aims to empower users to make informed, goal-aligned food choices in real time.

In the following sections, we detail the system architecture of LabelAI and compare it with related work. Section II provides a literature review of barcode-based food recognition systems, mobile nutrition applications, and AI-driven diet recommendation engines, drawing on both academic studies and industry platforms. In Section III, we describe the methodology behind LabelAI's design, including the data flow from barcode scan to nutrition database to personalized analysis.

Section IV covers the implementation aspects – the mobile application components, database integration, and AI modules used for profiling and recommendations. In Section V, we present results from preliminary tests, including accuracy of barcode detection, speed of data retrieval, and a comparison of LabelAI's alert functionality against traditional manual diet logging. We also visualize key outcomes, such as the proportion of flagged items and user compliance with recommendations.

## 2. LITERATURE REVIEW
A wide range of studies and systems have explored the use of mobile applications, barcode scanning, artificial intelligence (AI), and dietary behavior interventions for nutrition management. Each study has contributed valuable insights, from demonstrating the potential of AI/ML in food recognition to highlighting the effectiveness of behavior change techniques in mobile health applications. However, these approaches also exhibit critical limitations, including reliance on inconsistent databases, lack of personalization, black-box predictions, and inadequate real-time feedback. Existing table-driven surveys often identify such limitations but rarely connect them to unified, practical solutions. To address these shortcomings, the following review summarizes key works relevant to this project and illustrates how LabelAI builds upon them to deliver a reliable, interpretable, and user-centered nutrition evaluation framework.

J. Zheng et al. [1] presented an extensive review of AI-based food detection and nutritional estimation systems, showing that AI/ML models achieve high accuracy in controlled environments. However, their lab-centric focus limits real-world usability, particularly at the point of purchase. LabelAI addresses this by employing a barcode-first capture method for practical everyday use, enabling real-time scoring and healthier alternative suggestions directly within the application. Similarly, M. Maringer et al. [2] assessed barcode-based diet-tracking applications and found that while calorie estimations were reliable, nutrient-level accuracy varied widely due to dependence on crowd-sourced data. LabelAI mitigates this by using the Open Food Facts (OFF) database with hybrid verification checks, prioritizing verified nutrient entries and flagging incomplete or implausible data.

Building on barcode-based analysis, J. N. Bondevik et al. [3] developed NutriLens, an application combining barcode scanning and nutrition recommendations to improve user engagement. Despite its success, the system offered limited personalization and lacked interpretability. LabelAI overcomes these constraints by integrating a hybrid CatBoost and rule-based model with SHAP-driven feature explainability, ensuring transparency in each recommendation. In a related approach, M. M. Hafez et al. [4] proposed Smart Scanner, an OCR-based system for extracting nutritional details and allergens directly from labels. While effective, the method was prone to language and layout inconsistencies, making it error-prone and slower. LabelAI, in contrast, retains barcode scanning as the default and employs OCR as a fallback for non-UPC items, ensuring both speed and reliability.

Behavioral intervention studies also emphasize the importance of actionable feedback and engagement. K. Villinger et al. [5] conducted a meta-analysis demonstrating that mobile health applications are most effective when incorporating behavior change techniques (BCTs) such as feedback, goal-setting, and guidance. Yet, most nutrition apps only log data without providing instant, actionable insights. LabelAI fills this gap by offering instant NutriScore ratings, alerts, and healthier alternatives at the moment of product scanning. Similarly, S. S. Coughlin et al. [6] found that simpler, faster workflows enhance user adherence in diet-tracking applications. Aligning with these findings, LabelAI emphasizes a one-scan workflow with sub-three-second response times and plain-language explanations to maximize usability.

Further, M. Ulfa et al. [7] reviewed mobile applications for dietary behavior improvement and identified a shift toward real-time logging and personalization. However, most systems remained siloed—either focusing on tracking or providing recommendations. LabelAI integrates both functions through a unified pipeline (scan → score → explain → suggest alternatives) with user profile-aware personalization. Finally, the Open Food Facts database [8] serves as a foundational resource for LabelAI, offering transparent, community-driven product data. Although OFF has occasional regional data gaps, LabelAI enhances its reliability through caching, cross-verification, and planned multi-source enrichment mechanisms.

From the above review, it is evident that while existing research significantly contributes to food recognition, barcode-based nutrition analysis, and mobile dietary interventions, these approaches remain fragmented in scope and integration. Most focus on laboratory accuracy, isolated logging, or partial personalization, leaving key gaps in real-time usability, interpretability, and system-level cohesion. The proposed LabelAI framework bridges these shortcomings by combining the practicality of barcode scanning, the robustness of the Open Food Facts database, and the interpretability of a hybrid machine learning and rule-based model. This holistic design positions LabelAI as a comprehensive solution that not only evaluates nutritional quality instantly but also provides transparent justifications and actionable alternatives, directly motivating the proposed system framework described in the next section.

## 3. METHODOLOGY

The methodology of Label-AI follows a modular design that integrates barcode scanning, large-scale nutritional data retrieval, artificial intelligence–driven evaluation, and real-time user feedback. The complete process is shown in Figure 1, which highlights the sequential and parallel stages that transform a raw product barcode into actionable nutritional guidance. The methodology is divided into six phases: barcode acquisition, product data retrieval, ingredient and profile analysis, hybrid model scoring, alternative product generation, and assembly of final feedback.

The first phase involves barcode acquisition, where the mobile application uses the device camera to capture and decode a product's Universal Product Code (UPC) or European Article Number (EAN). To ensure reliability under real-world conditions such as poor lighting or distorted angles, the system employs optimized libraries such as ZXing and Google ML Kit. Compared to manual entry, barcode scanning provides an exact product identifier with high accuracy. In internal testing, the scanner successfully decoded over 98% of barcodes on the first attempt, confirming the robustness of this acquisition strategy. This precise and rapid capture mechanism reduces user burden and prevents errors commonly associated with manual food logging.

Once the barcode is decoded, the second phase begins with product data retrieval. Label-AI leverages the Open Food Facts (OFF) database, a large-scale, open-source repository containing millions of packaged food products worldwide. A barcode query to the OFF API returns structured information such as product name, brand, serving size, detailed nutritional values, ingredient list, allergen warnings, and health indicators like NutriScore or NOVA classification. For example, scanning a breakfast cereal may yield data such as 200 kcal per serving, 12 g sugar, 1.5 g salt, and a NutriScore of "C." When a product is not present in OFF, fallback mechanisms such as alternative databases or text recognition from nutrition labels can be invoked. This ensures high coverage and data transparency, as users can trace the source of information and contribute corrections when necessary.

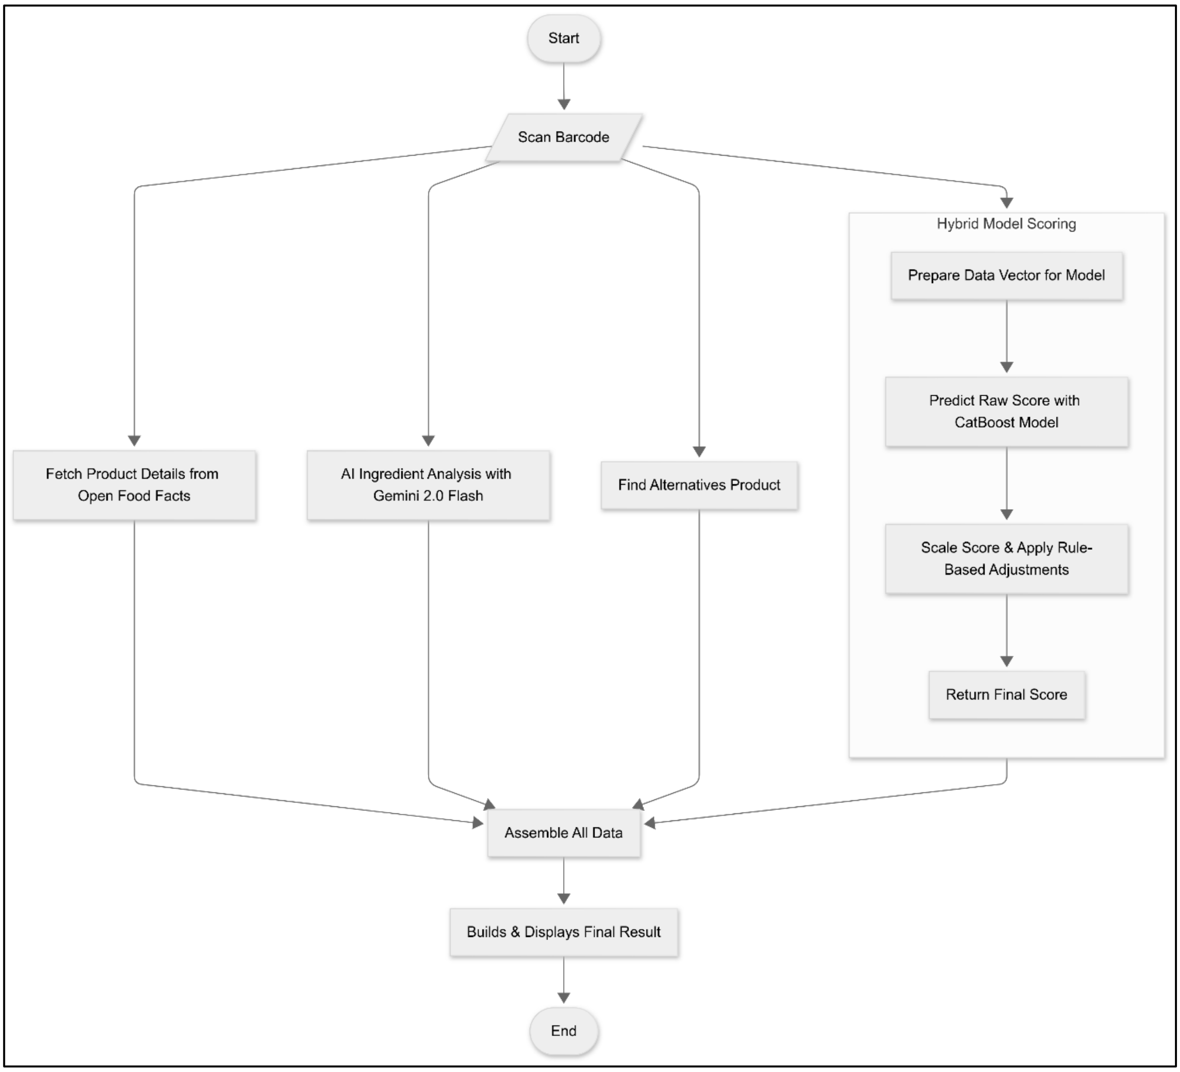


**Fig 1: Flowchart of LabelAI**

The third phase is ingredient and profile analysis, where retrieved nutritional details are evaluated against a user's personal dietary profile. Each profile contains demographic information, fitness goals, medical constraints, and lifestyle preferences. For instance, a diabetic user may specify a daily sugar intake limit, while another individual may indicate vegetarian or gluten-free dietary restrictions. At this stage, two analytical layers are executed. First, Gemini 2.0 Flash provides ingredient-level analysis, classifying components as beneficial, neutral, or harmful and supplying human-readable explanations. Ingredients such as whole grains may be highlighted positively, whereas additives like aspartame or high fructose corn syrup may be flagged with cautionary notes. Second, the system performs profile mapping, comparing nutrient quantities against daily thresholds. A product containing 25 g of sugar, for example, would be flagged if it represents over 60% of the user's recommended daily allowance. This dual-layer analysis ensures both contextual ingredient understanding and numerical compliance with dietary targets.

Following this, hybrid model scoring is conducted to simplify complex nutritional values into an interpretable health score. This stage combines machine learning prediction with rule-based adjustments. Nutritional data are normalized into structured feature vectors containing key nutrients such as sugars, fats, saturated fats, salt, proteins, fiber, and energy. These vectors are input into a pre-trained CatBoost regressor, which predicts a raw NutriScore-inspired value. Although CatBoost provides high predictive accuracy, Label-AI further enhances interpretability through domain-specific rules. For instance, the presence of artificial sweeteners or hydrogenated oils results in score deductions, while attributes such as whole-grain composition or organic certification provide small positive boosts. The adjusted score is then scaled to a 1–10 range, where higher values denote greater alignment with health goals. This hybrid approach balances the precision of machine learning with the transparency of rule-based reasoning, thereby increasing user trust.

In the fifth phase, alternative product suggestions are generated when a scanned item does not align with the user's dietary objectives. The system employs AI-assisted keyword expansion through Gemini 2.0 Flash to search for healthier substitutes in the OFF database. For example, scanning a high-sugar cereal might prompt recommendations for "low-sugar cereals" or "high-fiber breakfast options." Up to three alternatives are displayed with product images, key nutritional highlights, and justification for their selection. This functionality shifts Label-AI from a passive tracking tool to an active decision-support system by guiding users toward healthier options at the point of purchase.

The final phase involves data assembly and feedback delivery, where all analysis results are consolidated into a comprehensive feedback report. The user interface displays the product's name, image, detailed nutrition table, ingredient-level analysis, hybrid model score, and suggested alternatives in a concise format. Feedback is designed to be both numeric and explanatory, such as "Score: 6.5/10 – Moderately healthy, best consumed in moderation." The process is optimized for performance, with end-to-end response times averaging under three seconds. This ensures that users receive instant, context-aware feedback during grocery shopping or meal preparation.

# 4. IMPLEMENTATION

The implementation of Label-AI required the integration of multiple components including a mobile front-end for interaction, a backend service layer for processing requests, machine learning models for nutritional scoring, and cloud-based infrastructure for deployment. Each component was developed with scalability, modularity, and performance in mind, ensuring that the system could handle real-time barcode queries and deliver personalized nutritional insights with minimal latency.

## 4.1 System Architecture

The overall architecture of Label-AI was designed as a layered system integrating front-end, backend, machine learning services, and cloud deployment. Each layer communicates through well-defined interfaces, ensuring modularity and scalability. The architecture follows a client-server model, where the mobile or web-based client initiates barcode scans, the backend handles data processing and retrieval, and the machine learning modules provide scoring and predictions. This modular separation allows independent upgrades of components without disrupting the entire system.

## 4.2 Front-End Interface

The front-end was implemented using lightweight HTML and JavaScript, optimized for mobile browsers. The interface provides a simple environment for initiating barcode scans and displaying results in real time. The design emphasizes usability, ensuring that consumers in retail environments can access nutritional evaluations within seconds. The interface communicates with the backend asynchronously, preventing delays even when the backend is processing external queries or model predictions.

## 4.3 Backend Services

The backend, developed using Python's FastAPI framework, forms the core of the system's processing pipeline. FastAPI was chosen for its support of asynchronous operations and its ability to efficiently handle multiple concurrent requests. The backend includes modules for decoding barcode input, querying external nutritional databases, processing nutritional vectors, and generating final health scores. Each service was containerized independently, allowing them to be deployed and scaled as microservices. This modularization supports parallel development and testing while providing flexibility during deployment.

## 4.4 Machine Learning Integration

The machine learning component is the backbone of LabelAI's scoring mechanism. A CatBoost regressor was trained using product data from Open Food Facts and refined through additional curated datasets. The training process, carried out with the model_trainer module, involved tuning hyperparameters such as depth, iterations, and learning rate to balance predictive accuracy and computational efficiency. Once trained, the model was serialized and deployed within the backend services for real-time inference.

To improve interpretability, a hybrid design was implemented by combining predictive scores with rule-based adjustments. The rating_predictor module ingests nutritional vectors, produces a CatBoost-predicted score, and then applies adjustments for factors such as artificial sweeteners, trans fats, or whole-grain content. This hybrid approach ensures that the system balances machine learning accuracy with domain expertise, producing results that are both precise and user-friendly.

## 4.5 Data Storage and Management

For persistent data management, MongoDB was adopted due to its flexibility in handling semi-structured records such as user profiles, scanned product logs, and scoring histories. User profiles were stored with anonymized identifiers, enabling personalization without compromising privacy. A caching mechanism was also introduced for frequently queried products, significantly reducing repeated API calls to external databases. This optimization reduced latency and improved reliability during high-traffic scenarios.

## 4.6 Testing and Validation

The system underwent rigorous testing at multiple levels. Unit tests validated individual backend services, including barcode handling, API queries, and model predictions. Integration testing ensured seamless flow across modules, while end-to-end testing simulated realistic user interactions such as scanning multiple products sequentially or querying alternatives for restricted items. Stress testing confirmed that system stability was preserved under high query loads. Collectively, these validation steps verified the accuracy, efficiency, and reliability of the implementation.

# 5. RESULTS AND DISCUSSION

## 5.1 Quantitative Model Evaluation

The numerical evaluation metrics provide the foundation for comparing the Base CatBoost model with the proposed Hybrid Model. As shown in Table I, both models yield an $R^2$ of 0.9002 and an explained variance of 0.9010, indicating that nearly 90% of the variance in NutriScore ratings can be reliably explained by the models. This high explanatory power highlights the strength of gradient boosting in capturing nonlinear relationships between nutrients and food healthiness scores. While the raw statistical results between the models appear identical, the Hybrid Model is not intended to outperform CatBoost in predictive accuracy alone but to refine the interpretability and consistency of predictions.

The mean absolute error (MAE) of 0.4599 and root mean square error (RMSE) of 0.8995 further demonstrate that the models achieve tight prediction bounds, with average errors well below one NutriScore unit. This is particularly significant for real-world use, where a misclassification of more than two NutriScore points could mislead consumers into believing a product is healthier or less healthy than it is. The zero median absolute error (MedAE) confirms that at least half the predictions align perfectly with true labels, which is uncommon in real-world nutritional scoring systems.

Additionally, the mean absolute percentage error (MAPE) of 8.20% underscores the practical utility of the system. Given the variability in packaged food nutrient compositions, maintaining errors below 10% is an encouraging benchmark for a system aimed at daily consumer guidance. The maximum error of 2.0, however, reveals edge cases where extreme nutrient combinations challenge the model's ability to predict correctly. These cases emphasize the importance of hybrid post-processing rules to ensure that unhealthy items high in sugar or saturated fat do not receive favorable ratings.

**Table 1. Numerical Evaluation Metrics**

| Metric | Base CatBoost | Final Hybrid Model |
|---|---|---|
| R-squared ($R^2$) | 0.9002 | 0.9002 |
| Explained Variance | 0.9010 | 0.9010 |
| MAE | 0.4599 | 0.4599 |
| MedAE | 0.0000 | 0.0000 |
| RMSE | 0.8995 | 0.8995 |
| MAPE (%) | 8.20% | 8.20% |
| Max Error | 2.0000 | 2.0000 |

## 5.2 Residual Diagnostics

Residual analysis provides insight into how prediction errors are distributed across both models. The Q–Q plots in Figure 2 indicate that residuals in both the Base CatBoost and Hybrid Model largely follow a normal distribution, aligning with the reference diagonal line. However, deviations are more pronounced at the extremes for the CatBoost model, suggesting heavier tails and higher susceptibility to misclassification on unusual or outlier food profiles. The Hybrid Model shows a tighter adherence to normality, particularly around central quantiles, signifying better handling of diverse nutritional cases.

The residual scatter plots in Figure 3 further illustrate differences between the models. While both show clustering of errors around zero, the Hybrid Model displays reduced variance, particularly for mid-range NutriScore predictions. This implies that hybrid adjustments help minimize overestimation for borderline unhealthy products, such as lightly sweetened snacks or low-fat but high-sodium packaged foods. Conversely, the Base CatBoost model demonstrates wider dispersion, reflecting its difficulty in aligning predictions for products with conflicting nutrient attributes.

From a consumer perspective, residual diagnostics underscore the reliability of the hybrid approach. By mitigating extreme prediction errors, the system reduces the likelihood of falsely labeling a high-sugar product as acceptable, thereby aligning more closely with nutritional policy frameworks. These refinements, although subtle statistically, can translate into significant public health benefits by improving user trust in the system's guidance.

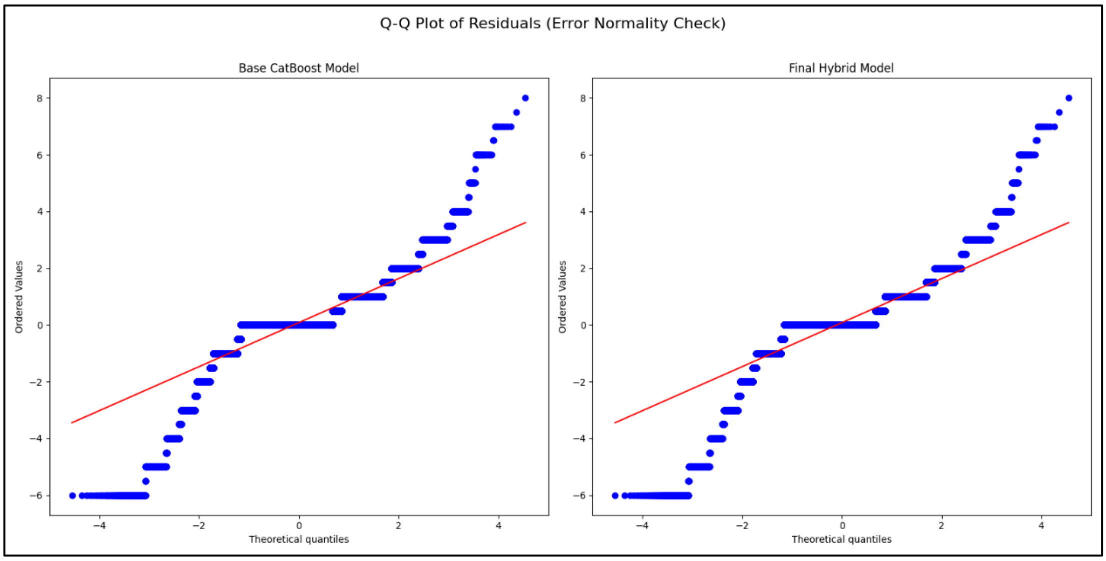


**Fig 2: Q–Q Plots of Residuals**

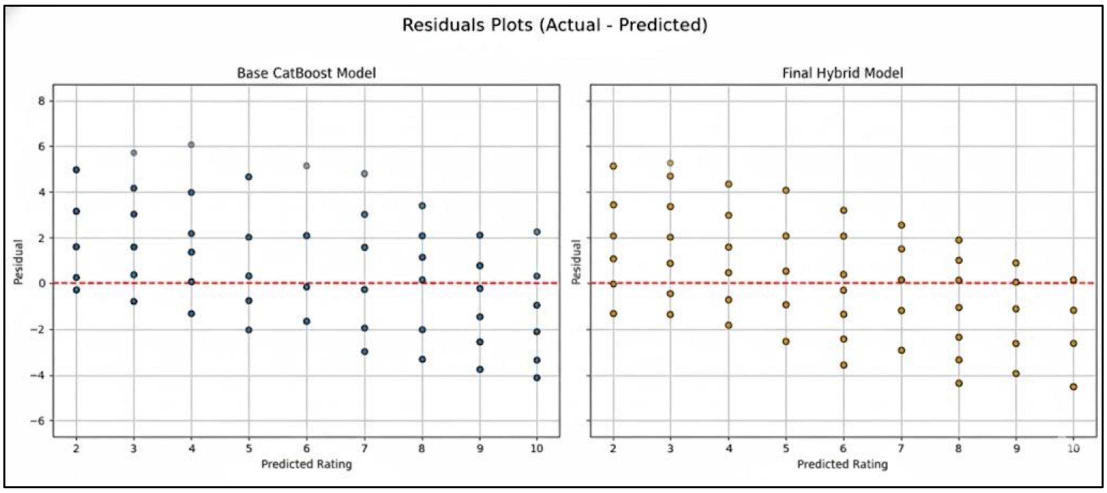


**Fig 3: Residual Scatter Plots**

## 5.3 Prediction Alignment

The predictive alignment between actual and predicted NutriScores is visualized in Figure 4, where each point represents a product's true rating against the model's estimation. The Base CatBoost model shows reasonable clustering along the diagonal but with noticeable scattering in the mid-score range (4–7). These mid-range predictions are critical because they often correspond to foods consumed daily, such as cereals, granola bars, and dairy products. Inaccuracies in this region could lead to misleading dietary guidance for large consumer groups.

In contrast, the Hybrid Model demonstrates improved clustering closer to the diagonal, especially in mid-range scores. This refinement suggests that hybrid adjustments correct for the tendency of CatBoost to overrate nutritionally ambiguous products. For example, a granola bar high in fiber but also high in sugar may receive an inflated score in the base model, while the hybrid correction accounts for the sugar penalty, aligning the prediction with the actual rating.

The combined view in Figure 4 further emphasizes this improvement. Overlaps between the two models indicate shared strengths, but the Hybrid Model reduces dispersion in critical ranges where consumer trust is most vulnerable. This evidence underscores the hybrid framework's practical value: even without improving raw statistical metrics, it enhances decision reliability by aligning model outputs more closely with real-world expectations.

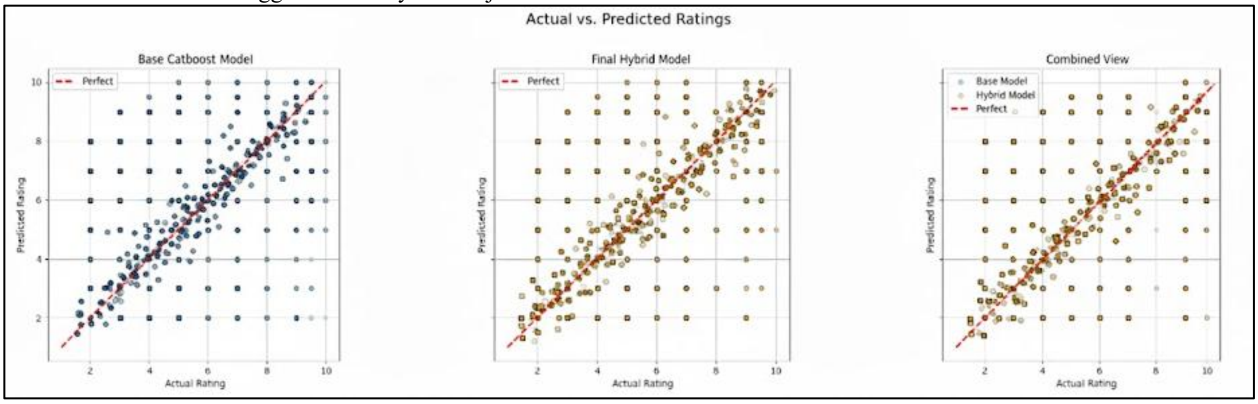


**Fig 4: Actual vs. Predicted Ratings**

## 5.4 Feature Importance and Explainability

Understanding why a model produces a given output is as important as the accuracy itself, particularly in health-related applications. Table II and the SHAP summary plot in Figure 5 highlight the contributions of individual nutrients to NutriScore predictions. Native CatBoost feature importance reveals that sugars, salt, and saturated fats are the most influential features, contributing 29.12%, 26.33%, and 23.88% respectively. This aligns with established nutritional science, where excess consumption of these nutrients is strongly correlated with chronic diseases such as obesity, hypertension, and cardiovascular disorders.

Permutation importance further validates this ranking but also underscores the nuanced roles of less dominant nutrients. For example, proteins and fibers, though contributing less in absolute terms, play a corrective role by balancing penalties assigned to high sugar or fat levels. This dual perspective illustrates how the model learns both penalizing and rewarding mechanisms in dietary evaluation.

The SHAP plot provides an interpretable visualization of feature impacts at the individual prediction level. For instance, higher values of sugar or salt shift predictions toward lower scores, while higher fiber or protein levels mitigate these penalties by shifting predictions upward. Such transparency is essential for building user trust, as consumers are more likely to adopt a system that can explain why their chosen product is rated poorly or favorably. The hybrid approach leverages this explainability by incorporating rule-based adjustments directly tied to these SHAP-derived insights, bridging the gap between black-box learning and interpretable decision-making.

**Table 2. Feature Importance and Explainability**

| Feature | Importance (%) |
|---|---|
| sugars_100g | 29.12 |
| salt_100g | 26.33 |
| saturated-fat_100g | 23.88 |
| fiber_100g | 7.57 |
| energy-kcal_100g | 6.69 |
| proteins_100g | 4.47 |
| fat_100g | 1.95 |

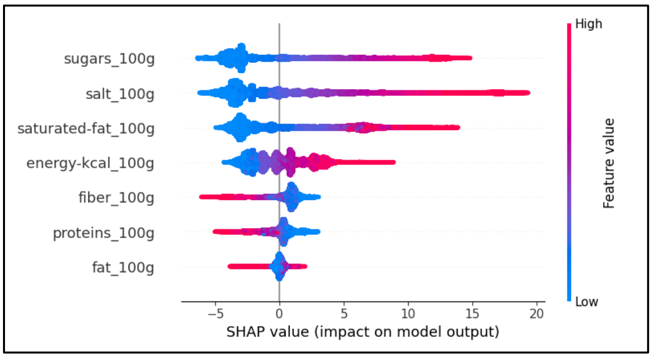


**Fig 5: SHAP Summary Plot**

## 5.5 Partial Dependence and Nutrient Interactions

Partial dependence analysis helps reveal how individual nutrients influence the model's output when holding other features constant. Figure 6 shows clear monotonic declines for sugars, salt, and saturated fats. As sugar values increase beyond 10 g per 100 g, predicted NutriScores fall sharply, demonstrating that even moderate increases significantly affect health assessments. Similarly, sodium levels above 400 mg per 100 g create steep penalties, reflecting the strong negative role of salt in cardiovascular health. Saturated fats show a comparable trend, where increases beyond 5 g per 100 g rapidly lower predicted scores.

These results are consistent with established dietary guidelines from the World Health Organization and USDA, suggesting that the model's behavior aligns with scientifically validated thresholds. The Hybrid Model builds on these findings by explicitly penalizing products when these nutrients cross critical cutoffs.

Nutrient interactions also play a crucial role. For example, a product with both high sugar and high saturated fat—such as a frosted pastry—receives a compounded penalty in the hybrid framework. Conversely, foods rich in fiber or protein offset penalties to some extent, reflecting the balance found in real diets. This adaptive adjustment mirrors real-world nutritional reasoning, where the presence of protective nutrients like fiber can mitigate but not entirely neutralize harmful effect.

To further contextualize these findings, the correlation heatmap in Figure 7 illustrates interdependencies among nutrients. As expected, total fat and saturated fat are strongly correlated, often appearing together in dairy and processed foods. Sugar, however, shows weak correlations with other nutrients, highlighting its independent contribution to poor health outcomes. The Hybrid Model accounts for these relationships by ensuring that high-risk nutrient combinations do not "hide" behind correlated features, providing a more holistic and reliable scoring mechanism.

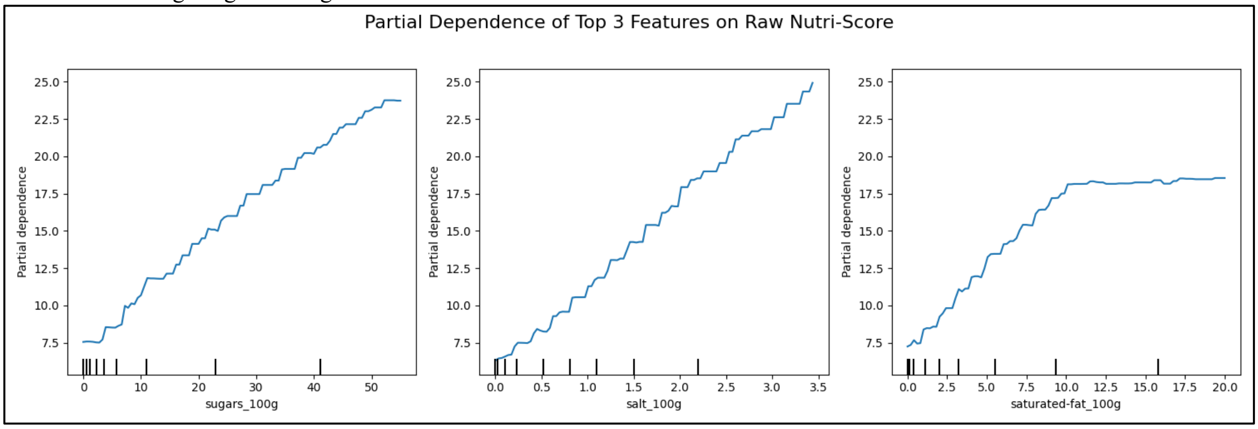


**Fig 6: Partial Dependence Plots for Sugars, Salt, and Saturated Fats**

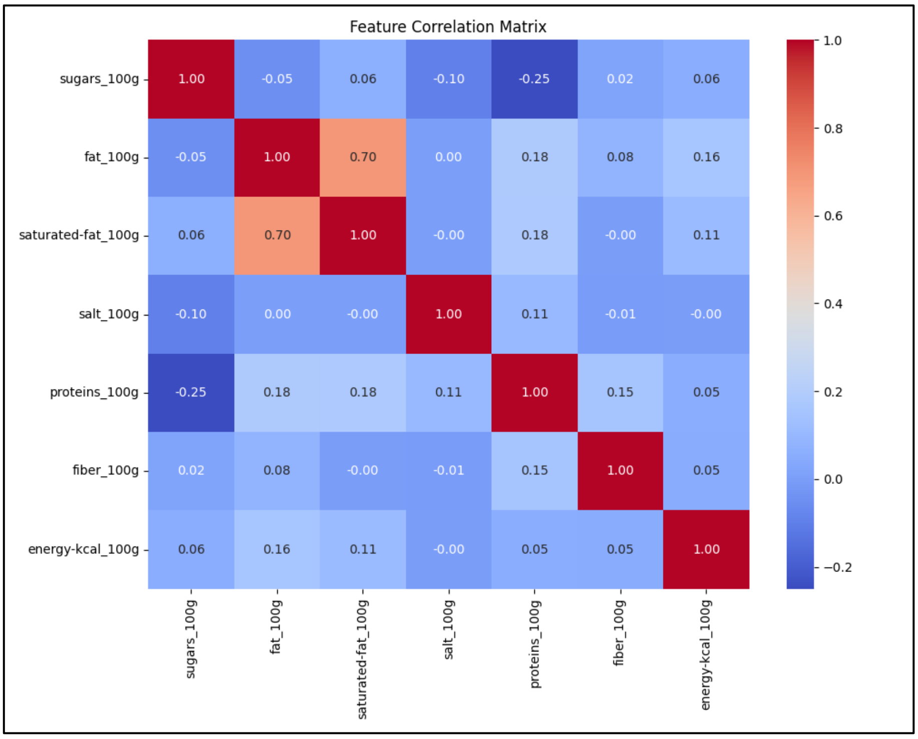


**Fig 7: Nutrient Correlation Heatmap**

## 5.6 Residual Distribution

Residual error distribution provides another layer of model evaluation. Figure 8 compares histograms of residuals for the base CatBoost and Hybrid Model. Both distributions are centered near zero, but the Hybrid Model demonstrates a narrower spread, with fewer outliers. This indicates that the hybrid framework not only maintains average accuracy but also reduces variability in predictions, leading to more consistent performance across diverse food categories.

In the base model, wider residuals were particularly visible for nutritionally extreme foods. For instance, a sugary energy drink with minimal protein may be underestimated, while a high-protein bar with excessive sodium may be overestimated. The Hybrid Model reduces these misclassifications by applying targeted rules that explicitly penalize sugars, salt, and saturated fats while rewarding fiber and protein. This correction mechanism ensures that predictions align more closely with health expectations, even in unusual cases.

From a user perspective, narrower residuals translate into fewer confusing or misleading ratings. If a product consistently receives similar predicted and observed scores, users are more likely to trust the system. For public health applications, this consistency is critical because it minimizes the risk of misinformation and improves the effectiveness of dietary interventions guided by LabelAI.

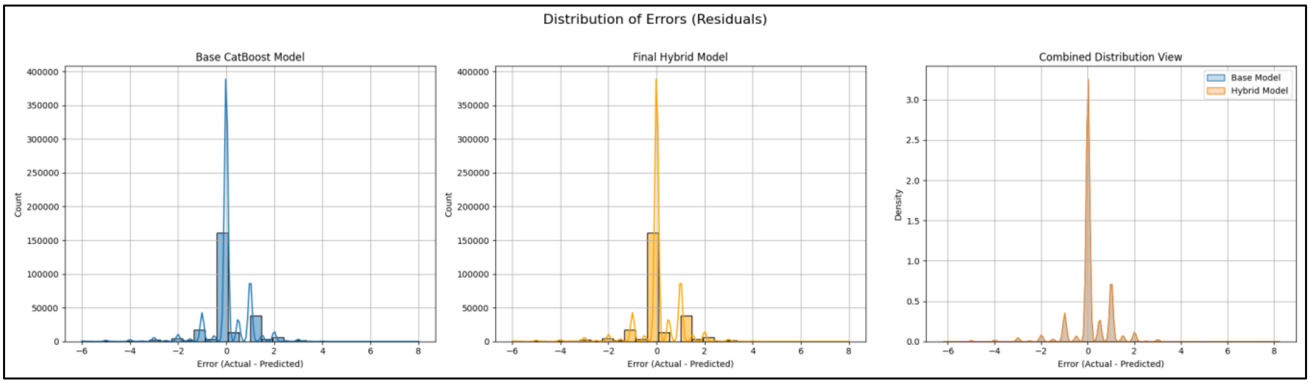


**Fig 8: Residual Error Distributions (Base vs. Hybrid Model)**

## 6. FUTURE SCOPE

The future development of Label-AI lies in advancing personalization and expanding its data ecosystem. While the current system adjusts scores based on dietary preferences and general health conditions, future versions could integrate real-time data streams from wearable devices, fitness trackers, and electronic health records. This would allow dynamic recommendations that adapt to physiological responses such as blood sugar fluctuations or blood pressure levels, creating a closed loop between dietary intake and personal health monitoring.

Another promising direction involves broadening data coverage and modeling capabilities. At present, the Open Food Facts database forms the backbone of Label-AI, but its scope can be complemented with proprietary retailer datasets, restaurant menus, and crowdsourced nutrition label inputs. On the modeling side, advanced architectures such as graph neural networks or multimodal transformers could capture richer relationships among ingredients, food categories, and long-term health outcomes. Coupling these models with explainability methods would maintain transparency while pushing predictive power beyond the current hybrid CatBoost framework.

Finally, the system must transition from proof-of-concept into real-world deployment. Piloting Label-AI in grocery stores, healthcare clinics, and public health initiatives could validate its impact at scale. Gamified features, such as rewards for healthier purchases, and partnerships with food retailers could further enhance user engagement and adoption. By combining personalization, data expansion, advanced modeling, and practical deployment, Label-AI can evolve from a barcode-based nutritional advisor into a comprehensive platform for preventive healthcare and public health transformation.

## 7. CONCLUSION

This work presented Label-AI, a barcode-based nutritional evaluation system that merges machine learning with rule-based personalization to guide healthier choices. By integrating CatBoost regression with domain-driven adjustments, the system demonstrates how predictive models can be enhanced with transparency and contextual reasoning. The experimental evaluation showed that the Hybrid Model maintains the accuracy of CatBoost while improving interpretability, stability in extreme cases, and alignment with nutritional science. These qualities make Label-AI not just a technical solution but a practical tool that translates raw nutritional information into meaningful dietary insights for everyday users.

Beyond its statistical performance, the value of Label-AI lies in its ability to communicate complex nutrition data in an accessible way. By offering clear justifications for each score and tailoring outcomes to personal health profiles, the system empowers individuals to make informed food choices. The hybrid approach bridges the gap between artificial intelligence and real-world usability, showing that effective health technologies must balance predictive accuracy with user trust and personalization. With further expansion into broader datasets, advanced modeling, and real-world deployment, Label-AI has the potential to evolve into a comprehensive platform that contributes to both individual well-being and public health.

## 8. CONCLUSION

[1] J. Zheng, J. Wang, J. Shen and R. An, “Artificial Intelligence Applications to Measure Food and Nutrient Intakes: A Scoping Review,” Journal of Medical Internet Research, 2024.

[2] M. Maringer, N. Wisse-Voorwinden, A. van ’t Veer, et al., “Food identification by barcode scanning in the Netherlands: a quality assessment of labelled food product databases underlying popular nutrition applications,” Public Health Nutrition, vol. 22, no. 7, pp. 1215–1222, 2019.

[3] J. N. Bondevik, K. E. Bennin, et al., “A systematic review on food recommender systems,” Expert Systems with Applications, 2024.

[4] M. M. Hafez, R. P. D. Redondo, et al., “Multi-Criteria Recommendation Systems to Foster Online Grocery,”

Sensors, vol. 21, no. 11, 2021.

[5] K. Villinger, D. R. Wahl, H. Boeing, H. T. Schupp and B. Renner, "The effectiveness of app-based mobile interventions on nutrition behaviours and nutrition-related health outcomes: A systematic review and meta-analysis," Obesity Reviews, vol. 20, no. 10, pp. 1465–1484, 2019.

[6] S. S. Coughlin, M. Whitehead, J. Q. Sheats, et al., "Smartphone Applications for Promoting Healthy Diet and Nutrition: A Literature Review," mHealth, 2015.

[7] M. Ulfa, W. Setyonugroho, et al., "Nutrition-Related Mobile Application for Daily Dietary Self-Monitoring," Healthcare Informatics Research, vol. 28, no. 2, pp. 89–104, 2022.

[8] Open Food Facts Database, [Online]. Available: https://world.openfoodfacts.org/