

ABSTRACT

This study introduces and validates a recommendation framework that uses a dual-classifier XGBoost model to assign personalized physical exercise and diet templates based on a user's fitness goals. The framework addresses the limitations of existing systems, which often rely on simplistic rules or generic recommendations, by using a machine learning approach trained on a robust dataset derived from evidence-based templates.

Performance evaluation showed that the real-world XGBoost models achieved an average test accuracy of 74.4% (82.1% for diet and 66.7% for physical exercise), a result that proved to be practical for deployment. A comparative analysis against a theoretical "Perfect" model revealed a significant performance gap of up to 31.8%. This gap demonstrates that the primary limitation is not the algorithm but the inherent ambiguity, noise, and severe class imbalance of the training data. The model's superior performance on the diet task was attributed to a clearer feature hierarchy, where BMI-related variables provided more deterministic decision boundaries than the complex, multi-feature challenge of predicting physical exercise templates.

Despite a marginal accuracy difference from a comparable Random Forest model, XGBoost was selected for its superior computational efficiency, with a model size 91.81% smaller and inference time 79.5x faster. This makes it ideal for a responsive web application. User acceptance, validated through a survey and a case study, was high (4.19/5 overall satisfaction), indicating that the system successfully translates complex machine learning predictions into a valuable and user-friendly experience. This research confirms the potential of machine learning for evidence-based fitness recommendations and establishes a foundation for future development focused on addressing data quality to enhance performance.

Keywords: *XGBoost, fitness recommendation, template classification, website application.*