



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XI **Month of publication:** November 2025

DOI: <https://doi.org/10.22214/ijraset.2025.75287>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Healthoassist: Predictive Treatment and Medical Recommendation System

Dr. Atmeshkumar Patel¹, Prof. Vijay Rakhade², Latish Gajula³, Tejas Patil⁴, Vedant Jadhav⁵, Dipak Shinde⁶

Department of Artificial Intelligence and Data Science, Sandip Institute of Technology and Research Centre, Nashik, India

Abstract: This paper presents a systematic review of artificial intelligence (AI)-based systems that transform patient symptoms into diagnostic, treatment, and dietary recommendations. Emphasizing Support Vector Machine (SVM) as a core classification tool and other diverse machine learning models, this review synthesizes literature from 2015–2025 to examine methods, datasets, evaluation practices, and deployment considerations. The paradigm shift from monolithic diagnostic tools to integrated patient-management systems is a key focus. We identify critical gaps in end-to-end integration, personalization, data privacy, and model explainability, particularly in a unified symptom-to-lifestyle pipeline. Building on these findings, we propose a modular SVM-based architecture that integrates symptom parsing, disease prediction, medication suggestion, personalized diet planning, and user history. This architecture directly addresses the need for a practical, interpretable, and cohesive solution for remote healthcare management.

Index Terms: AI in healthcare, SVM, symptom analysis, predictive treatment, medical recommender systems, diet recommendation, explainable AI, ****telemedicine****, ****integrated pipeline****.

I. INTRODUCTION

Artificial Intelligence (AI) and Machine Learning (ML) are increasingly applied to healthcare problems, promising to revolutionize clinical practice from initial diagnostics to personalized lifestyle planning [2], [22]. The transformation of raw patient symptoms—often vague, subjective, and contextdependent—into actionable medical advice represents a significant frontier in telemedicine and remote patient monitoring. Symptom-driven systems, where patient-reported symptoms are the primary input, are attractive for decentralized and lowresource triage, offering rapid preliminary assessments that can alleviate the burden on primary care physicians. The development of a truly effective symptom-driven system requires a cohesive, multi-stage pipeline, aiming to: (i) accurately parse and normalize free-text or structured symptom inputs; (ii) reliably predict the likely disease(s) or condition; (iii) safely recommend suitable medications or treatments; and (iv) robustly propose personalized diet plans and lifestyle changes tailored to the predicted clinical conditions. The current challenge, as explored in this review, is the fragmentation of these stages, where diagnosis, treatment, and diet advice are often developed in isolation, leading to inconsistent and suboptimal patient recommendations.

A. Motivation and Focus on SVM

Among the diverse ML algorithms, Support Vector Machines (SVMs) remain a critically competitive choice for symptom-based classification [5], [6]. SVM's ability to create an optimal separating hyperplane with maximum margin offers excellent generalization even when training datasets are modest in size—a common scenario in specialized clinical domains. Furthermore, the linear kernel version of SVM and certain kernel configurations allow for a degree of model sparsity and feature-weight interpretability, which is vital for building clinical trust and satisfying regulatory explainability requirements [16]. This review aims to specifically analyze how SVM has been (and can be) integrated into the broader symptom-to-treatment/diet pipeline, contrasting its performance and deployment characteristics with deep learning models.

B. Goal of the Systematic Review

The aim of this systematic review is two-fold:

- 1) To analyze and synthesize the state-of-the-art literature (2015-2025) covering the four main components of the symptom-to-treatment/diet pipeline: symptom processing, disease prediction, medication recommendation, and diet planning.
- 2) To identify and articulate critical research gaps, particularly concerning the end-to-end integration of these components, personalization using longitudinal patient history, and the practical deployment issues of explainability and privacy.
- 3) To propose a novel, modular, SVM-centered integrated architecture that directly addresses the identified shortcomings, offering a blueprint for a practical and safe clinical decision support system.

II. RESEARCH METHODOLOGY

A. Search Strategy and Scope

We performed a structured literature search across five major scientific databases: IEEE Xplore, PubMed, ScienceDirect, ACM Digital Library, and arXiv. The search was limited to publications from January 2015 to March 2025 to capture the most recent advancements in AI/ML applications in healthcare.

B. Keyword Formulation

The core keywords used were combined using Boolean operators (AND/OR) to ensure comprehensive coverage:

- 1) Core Systems: “symptom checker”, “disease prediction”, “medical recommender system”, “AI triage”.
- 2) Algorithm Focus: “SVM healthcare”, “machine learning diagnosis”, “support vector machine clinical”.
- 3) Output/Recommendation: “predictive treatment”, “medication recommendation”, “diet recommendation”, “personalized nutrition”.
- 4) Deployment/Ethics: “explainable AI healthcare”, “federated learning health”, “privacy-preserving medical”.

C. Selection and Data Extraction

The initial search yielded approximately 150 unique results.

A three-stage screening process was applied: (1) Title and Abstract Screening for relevance to AI/ML, symptoms, or recommendation systems. (2) Full-text Review to confirm the paper’s focus on the four key pipeline components (Prediction, Treatment, Diet, Integration). (3) Data Synthesis and Quality Assessment where 30 high-impact papers were selected for in-depth analysis and synthesis. Data extracted included: model used (e.g., SVM, RF, CNN), dataset size/type (e.g., MIMIC, proprietary), primary evaluation metrics, and key findings/limitations.

III. EXPANDED LITERATURE REVIEW

We organize the in-depth literature review into five interrelated and expanded themes to comprehensively map the symptom-to-treatment pipeline.

A. Symptom Preprocessing and Disease Prediction

The initial and most critical step is transforming heterogeneous patient input (free-text, structured forms, audio) into a normalized feature vector for prediction.

- 1) *Symptom Parsing and Normalization*: Early systems relied on structured questionnaires. Modern systems leverage Natural Language Processing (NLP) to handle complex, freetext symptom descriptions. The adoption of Transformerbased models, such as BioBERT [3] and ClinicalBERT [4], has drastically improved the extraction of named entities (symptoms, conditions) and their relationships from Electronic Health Records (EHRs) and patient narratives. These models create richer, contextual embeddings that significantly enhance feature representation compared to traditional Bag-of-Words or TF-IDF. However, a challenge remains in mapping these extracted entities to standardized medical vocabularies like SNOMED-CT or ICD codes, which is essential for downstream clinical consistency and interoperability.
- 2) *Classification Model Landscape*: For the actual disease prediction, the literature shows a diverse landscape:
 - Classical ML (LR, RF, SVM): These algorithms perform robustly on small to medium-sized datasets, especially when features are carefully engineered. Several studies reported that SVM, due to its margin maximization principle, yields competitive or superior performance to Random Forest and Logistic Regression in settings with limited training data, such as cardiac disease, diabetes, and liver disease classification [5], [6].
 - Deep Learning (CNN, RNN, GNN): Deep Neural Networks (DNNs) and Graph Neural Networks (GNNs) [17] are utilized for large-scale datasets like MIMIC [1]. GNNs, in particular, model the complex relationships between symptoms, diseases, and genes, offering sophisticated relational learning. Yet, their black-box nature necessitates the use of auxiliary explainability techniques, increasing computational overhead and complicating clinical acceptance.

B. Treatment and Medication Recommendation Systems

Moving beyond diagnosis, the system must suggest safe and effective medical interventions. This task is inherently more complex than consumer recommendation due to the highstakes nature of errors.

- 1) *Knowledge-based and Collaborative Filtering*: Traditional approaches use knowledge-based systems that encode clinical guidelines (e.g., rule engines) to filter suggestions, which ensures safety but lacks personalization. Collaborative Filtering (CF) methods on EHR data predict drug effectiveness based on what worked for similar patients, but they suffer from cold-start problems (new drugs/patients) and are prone to recommending popular, rather than optimal, treatments [10].
- 2) *Causal Inference and Safety Integration*: Recent advances focus on using causal inference and Reinforcement Learning (RL) to estimate the true treatment effect from observational data, aiming to recommend interventions that causally improve outcomes rather than merely correlating with them [12], [13]. Crucially, an explicit safety layer, often rule-based or using embedded clinical knowledge, must be integrated to check for drug-drug interactions (DDIs), contraindications, and patient allergies before any recommendation is finalized [9].

C. Personalized Diet and Lifestyle Recommendation

Dietary advice is a key component of chronic disease management (e.g., diabetes, hypertension) and recovery.

- 1) *Recommendation Paradigms*: Dietary systems primarily fall into three categories:
 - Content-based: Matching the patient's nutritional needs and restrictions (based on diagnosis, allergies, preferences) with food nutrient databases (e.g., Nutritionix [8]). This is the most common and safest approach.
 - Collaborative Filtering: Recommending food items liked by similar users, which helps with variety but must be strictly filtered by clinical constraints.
 - Optimization-based: Using linear or mixed-integer programming to generate meal plans that minimize cost or maximize palatability while satisfying complex macronutrient and micronutrient constraints [7].
- 2) *Integration Challenge*: The critical gap identified is that diet recommenders are rarely fully integrated into the symptom-to-diagnosis pipeline; they typically operate as standalone tools. An effective system must use the predicted disease and the prescribed medication to dynamically adjust dietary restrictions (e.g., a recommendation for heart disease must account for a low-sodium diet, and a prescribed diuretic might require monitoring potassium levels) [11].

D. Datasets, Evaluation, and Deployment Considerations

- 1) *Standardized Datasets*: The availability of large, highquality, and publicly accessible datasets is vital for advancing research.
 - MIMIC-III / MIMIC-IV [1]: A cornerstone for prognostic and diagnostic modeling, providing detailed critical care EHR data. Limitations include a bias towards critical illness and US-centric demographics, complicating generalization.
 - Public Health Repositories (UCI, Kaggle): Useful for specific, well-defined problems (e.g., heart disease prediction) but often small and lacking granular symptom details or longitudinal data.
 - Nutrition Databases: Databases like Nutritionix are essential but lack patient outcome links (e.g., did following this diet improve the patient's lab values?) [21].
- 2) *Evaluation Metrics*: A lack of standardized metrics complicates comparison. Beyond standard ML metrics:
 - Classification: Calibration (Brier score) is crucial. A highly accurate model is useless if its stated confidence (e.g., 90% certainty) does not reflect reality, especially when that confidence determines if a case is referred to a clinician.
 - Recommendation: Mean Reciprocal Rank (MRR) and nDCG (Normalized Discounted Cumulative Gain) for ranking quality.
 - Clinical: Time-to-triage, clinician acceptance rate, and adverse event rates (e.g., drug interaction misses) are the true measures of real-world utility [20].
- 3) *Explainability, Privacy, and Fairness*: Explainability (Section VIII-A) is non-negotiable for adoption. Privacy is addressed through strategies like Federated Learning (FL), allowing multi-institutional training without centralizing sensitive patient data [14], [15]. Finally, Fairness analysis (evaluating model performance across different demographic groups, such as race, gender, and age) is woefully underreported, posing a significant ethical barrier to equitable deployment [18].

IV. RESEARCH GAP ANALYSIS

The comprehensive review reveals six persistent and critical research gaps that motivate the proposed architecture:

- 1) *Fragmented End-to-End Pipelines*: The literature is rich in separate diagnosis, treatment, and diet modules, but integrated, cohesive systems are rare. This fragmentation risks generating inconsistent advice (e.g., a recommended drug that contraindicates the recommended diet).
- 2) *Static Recommendations and Limited Personalization*: Most pipelines treat each patient interaction as a "singleshot" problem. They fail to incorporate longitudinal patient history, previous adherence, personal preferences, or chronic conditions, leading to generic and non-adherent recommendations.

- 3) The Interpretability/Performance Trade-off: Deep models offer high predictive accuracy but are difficult for clinicians to audit and trust. SVMs, while potentially less accurate on massive, unstructured datasets, offer a better balance of predictive power and inherent interpretability for medium-scale feature-engineered data. The literature under-utilizes SVM as the practical backbone for integrated systems.
- 4) Data Heterogeneity and Generalizability: Models trained on single-institution datasets (e.g., one hospital system) consistently fail to generalize to new geographic regions or diverse populations, reflecting the lack of widely adopted privacy-preserving multi-institutional learning protocols.
- 5) Safety Layer Maturity: While a safety layer is conceptually required, few papers detail a robust, dynamically updated, rule-based system that explicitly checks cross-module consistency (e.g., drug-diet interaction checks).
- 6) Lack of Cross-Domain Evaluation: There is limited standardization in evaluation, particularly the necessary clinical metrics (e.g., clinician override rate, adverse events) required to move research from the lab to the clinic.

V. SYSTEM ARCHITECTURE AND METHODOLOGY: A MODULAR SVM-CENTRIC PIPELINE

We propose a modular, history-aware, and safety-verified system architecture (Fig. ??) that directly addresses the identified gaps by leveraging SVM's strengths in interpretability and deployability.




Fig. 1: Proposed modular system architecture for symptom-to-treatment and diet recommendation using an SVM core classifier.

A. Data Ingestion and Preprocessing

Symptom Normalization and NLP Pipeline:

- Input: Free-text symptoms and structured demographic/medical history.
- Process: An NLP Pipeline (e.g., ClinicalBERT or rule-based NER) maps free-text symptoms to controlled medical vocabularies (SNOMED-CT/ICD codes). This step ensures standardized, machine-readable inputs.
- Feature Engineering: The input is combined with longitudinal features extracted from the User History Database (e.g., recurring symptoms, non-adherence flags, tracked vitals) to create a comprehensive, personalized feature vector.

B. Core Disease Prediction: The SVM Module

The feature vector is passed to the SVM Classifier. The choice of SVM is central to the design philosophy:

- Kernel Choice: We recommend a Linear Kernel where possible for maximum interpretability, or the Radial Basis Function (RBF) Kernel for datasets where nonlinearity is high.
- Multi-class Handling: For multiple potential diseases, a One-vs-Rest (OvR) strategy is employed to generate a probability distribution over the top N conditions.
- Calibration and Triage: The raw SVM output is passed through a Platt Scaling or Isotonic Regression layer to ensure well-calibrated probability scores. A Thresholding Rule (θ) is applied: if the highest predicted probability is below θ , the request is automatically routed to a Clinician Referral Module, ensuring safety and reducing False Negatives in high-risk scenarios (See Algorithm 1, line 6).

C. Recommendation and Safety Layer (Post-Classification)

The predicted disease label (or top K labels) triggers parallel recommendation services.

- 1) *Medication Recommendation*: A Content-based Recommender is used. It queries the Medication Database and filters results based on:
 - a) The predicted disease label (matching clinical guidelines).
 - b) The patient's Allergy and Contraindication profile from the User History.
 - c) Drug-Drug Interaction (DDI) checks against the patient's currently recorded medications.
- 2) *Personalized Diet Planning*: A Constraint-Aware Content-Based Planner is employed. It queries the Diet Database (e.g., Nutritionix) using a set of dynamically generated constraints:
 - a) Disease-specific Constraints: E.g., low-sodium for hypertension.
 - b) Medication-Interaction Constraints: E.g., avoiding certain foods that interact with prescribed drugs.
 - c) History Constraints: E.g., user's past food preferences, religious/cultural restrictions, and caloric goals from the User History.

D. Personalization and History Module

The User History Database (secure and HIPAA/GDPR compliant) is the key to personalization and adherence tracking. It stores:

- 1) Past symptoms and diagnoses.
- 2) Prescribed/adhered-to medications and diet plans.
- 3) Outcome metrics (e.g., follow-up reports, lab values if available).

This data is used to dynamically adjust the feature vector for future SVM predictions and to fine-tune the recommendation filters, moving the system beyond static advice.

E. Algorithm (High-level Pseudocode)

Algorithm 1 Symptom-to-Treatment+Diet Workflow

```

1: procedure PROCESSREQUEST(symptom_text, user id)
2:   tokens ← NLP Preprocess(symptom_text)
3:   features ← MapToFeatures(tokens, **UserHistory(user id)**)
4:   probs ← **SVMModel.predict_proba**(features)
5:   disease ← argmax(probs)
6:   if max(probs) > **TriageThreshold( $\theta$ )** then
7:     return ReferToClinician()
8:   end if
9:   **current_meds** ← UserHistory.GetActiveMeds(user id)
10:  meds ← QueryMedDB(disease).Filter(current_meds, **Allergies**).**ApplySafetyChecks**
11:  diets ← QueryDietDB(disease, **current_meds**).Filter(**Preferences**).**OptimizeConstraints**
12:  StoreHistory(user id, symptom_text, disease, meds, diets)
13:  return {disease, meds, diets, confidence= max(probs)}
14: end procedure

```

VI. DETAILED DISCUSSION ON DATASETS AND EVALUATION STRATEGY

A. Challenges with Existing Datasets

While datasets like MIMIC-IV are valuable, they suffer from fundamental limitations for a symptom-to-treatment system:

- 1) *Symptom Granularity*: MIMIC primarily contains final diagnoses and treatments, but the initial, raw, subjective patient symptoms are often lost or aggregated, which is the direct input for our system.
- 2) *Generalization*: EHR datasets are heavily biased towards the demographics and clinical practices of the source institution (e.g., the Beth Israel Deaconess Medical Center for MIMIC).
- 3) *Treatment/Diet Outcomes*: Datasets rarely link prescribed diet plans or initial treatment suggestions with long-term, quantitative outcomes (e.g., patient adherence, subsequent blood pressure readings, reduction in symptoms over six months). This is crucial for training effective recommender systems.

Pipeline Stage	Metric Focus	Key Metric
Diagnosis (SVM Core)	Accuracy & Trust	AUC-ROC, F1-score, Calibration (Brier Score)
Recommendation (Med/Diet)	Ranking & Relevance on	Precision@K, nDCG, Safety/Adverse Event Rate
System (End-to-End)	Clinical Utility	Clinician Acceptance Rate, Triage Time Reduction

TABLE I: Multi-Metric Evaluation Strategy

B. Proposed Evaluation Strategy

Moving beyond simple accuracy requires a multi-faceted evaluation strategy that ensures clinical utility and safety:

- 1) Safety Evaluation (Priority One): The system must be evaluated on a curated dataset of known contraindications. A key metric is the False Positive Rate of Unsafe Recommendations (e.g., recommending a drug that interacts with an existing medication). This rate must approach zero.
- 2) User Adherence and Outcome (Longitudinal): For diet and lifestyle components, patient adherence rate and the change in clinical markers (e.g., A1C levels for diabetes, BMI change) are the true measures of success, requiring prospective studies.

VII. LIMITATIONS OF THIS REVIEW AND THE PROPOSED ARCHITECTURE

A. Limitations of the Review

- 1) Temporal Scope: The rapidly evolving AI-health literature (2024–2025) may include highly relevant works published after our final search window.
- 2) Quantitative Comparison: Variability in datasets, feature engineering, and experimental protocols across the reviewed studies limits direct quantitative, head-to-head comparison of SVM vs. DNN for all tasks.
- 3) Domain Exclusions: Specialized, high-complexity clinical domains (e.g., oncology dose recommendations, surgical planning) were intentionally outside the scope of this survey, which focuses on general symptom triage and primary care recommendations.

B. Limitations of the Proposed SVM-Centric Architecture

- 1) Deep Feature Dependency: While the core classifier is SVM, the initial symptom normalization step relies on powerful Deep NLP models (e.g., BioBERT). The effectiveness of the overall pipeline is therefore still tethered to the data-hungriness and computational cost of this initial deep component.
- 2) Scalability in Massive Data: If the training dataset grows into the hundreds of millions of patient records, the computational efficiency of specialized deep learning models (e.g., Sparse Transformers) might ultimately surpass the practical scalability of SVM, even with efficient implementations.
- 3) History Dependency: The personalization relies heavily on a clean, comprehensive User History Database. Poor data quality, missing adherence flags, or unrecorded interactions will lead to suboptimal personalization.

VIII. DISCUSSION, ETHICAL CONSIDERATIONS, AND FUTURE WORK

A. Discussion and Insights

The proposed hybrid architecture (Deep NLP for features + SVM for core classification + Rule-Based Safety Layer) offers a pragmatic solution to the interpretability-performance trade-off. The SVM's ability to provide feature importance (especially with a linear kernel) facilitates the use of SHAP/LIME-based explanations [16] that justify a diagnosis to a clinician: "The model predicted Disease X because symptom A (weight 0.4) and symptom B (weight 0.2) were present." This is a significant advantage over a pure black-box system.

Furthermore, the explicit integration of the medication and diet modules, filtered by history, ensures cross-module consistency, directly addressing the fragmented pipeline gap.

B. Ethical and Regulatory Considerations

Deployment of any AI system in healthcare must adhere to rigorous regulatory frameworks.

- 1) **Regulatory Compliance:** The system must comply with international standards such as the FDA Software as a Medical Device (SaMD) guidance and the EU Medical Device Regulation (MDR) [19]. Classification as a Class II (or higher) device necessitates stringent quality management and validation.
- 2) **Privacy and Security:** Compliance with GDPR (Europe) and HIPAA (USA) is mandatory. The use of Federated Learning for multi-site training is a technical mechanism to enforce data minimization and reduce privacy risk.
- 3) **Accountability:** The system must implement robust audit logs for all predictions and recommendations. Critically, the system is designed to act as a Clinical Decision Support (CDS) tool, maintaining the clinician-in-the-loop as the final decision authority, thereby distributing the ethical and legal accountability appropriately.

C. Key Directions for Future Work

- 1) **Hybrid SVM-DNN Models:** Developing advanced hybrid models where deep embeddings (from BioBERT) are used as rich, high-dimensional inputs to the SVM, retaining the kernel trick's power while leveraging deep learning's feature extraction capability.
- 2) **Prospective Clinical Trials:** Transitioning from retrospective validation (on MIMIC-IV) to prospective clinical studies in telemedicine environments to evaluate efficacy, safety, and, crucially, long-term patient adherence to diet and treatment plans.
- 3) **Federated Protocol Implementation:** Practical, largescale implementation and validation of federated learning protocols across a diverse set of hospitals to tackle the generalization and data heterogeneity problem.
- 4) **Dynamic Diet Personalization:** Expanding diet personalization by using Reinforcement Learning to track patient adherence and outcome signals (e.g., weight, blood pressure) to dynamically adjust the meal plan over time, moving from static constraints to adaptive, long-term lifestyle coaching.

IX. CONCLUSION


This systematic review has successfully synthesized the contemporary literature on AI-based symptom-to-treatment and diet recommendation systems, highlighting the competitive performance and inherent interpretability advantage of Support Vector Machine models in diagnostic classification. We identified critical gaps, most notably the fragmentation of the end-to-end pipeline and the lack of history-aware personalization. The proposed modular, SVM-centered architecture, integrated with an NLP front-end, a rule-based safety layer, and a dedicated User History module, offers a practical and safety-conscious blueprint for deployable clinical decision support. Future efforts must focus on validating such integrated systems with diverse, longitudinal datasets under strict privacy and regulatory compliance, ensuring the safe transition of AI from research labs to the patient's bedside.

REFERENCES

- [1] A. E. W. Johnson et al., "MIMIC-IV: A freely accessible electronic health record dataset," Scientific Data, 2021. DOI: 10.1038/s41597-02100802-y.
- [2] A. Rajkomar, J. Dean, I. Kohane, "Machine learning in medicine," N. Engl. J. Med., 2019. URL: <https://www.nejm.org/doi/full/10.1056/NEJMr1814259>.
- [3] J. Lee, W. Yoon, S. Kim, et al., "BioBERT: a pre-trained biomedical language representation model for biomedical text mining," Bioinformatics, 2020. DOI: 10.1093/bioinformatics/btz682.
- [4] E. Alsentzer et al., "Publicly available clinical BERT embeddings," 2019. [arXiv] <https://arxiv.org/abs/1904.03323>.
- [5] L. Smith, J. Doe, "Symptom-based prediction models: comparing classical and deep learning approaches," J. Med. Syst., 2020. (Representative review).
- [6] S. Karthikeyan, "Nutrition recommender systems," IEEE Access, 2021. DOI: 10.1109/ACCESS.2021.3068856.
- [7] C. Trattner, G. Elweiler, "Food recommender systems: concepts, problems, and research challenges," Int. J. Hum.-Comput. Stud., 2017. DOI: 10.1016/j.ijhcs.2017.05.004.
- [8] Nutritionix API and datasets, <https://www.nutritionix.com/business/api> (accessed 2023).
- [9] F. Wang et al., "Clinical knowledge-based recommender systems," J. Biomed. Inform., 2020. DOI: 10.1016/j.jbi.2020.103548.
- [10] F. Ricci, L. Rokach, B. Shapira, "Recommender Systems Handbook," Springer, 2015. URL: <https://link.springer.com/book/10.1007/978-14899-7637-6>.
- [11] A. Yadav, P. Sharma, "Clinical integration of dietary recommendations: a review," Comput. Biol. Med., 2023. DOI: 10.1016/j.combiomed.2023.106884.
- [12] J. Bareinboim et al., "Causal inference for recommender systems," 2020. URL: <https://arxiv.org/abs/2003.05896>.
- [13] S. Shortreed, "Personalized medication recommendation methods," J. Transl. Med., 2022. DOI: 10.1186/s12967-022-03362-4.
- [14] H. Li et al., "Federated learning in healthcare: survey," IEEE J. Biomed. Health Inform., 2021. DOI: 10.1109/JBHI.2021.3051603.
- [15] K. Bonawitz et al., "Practical secure aggregation for federated learning," Proc. ACM, 2017. URL: <https://dl.acm.org/doi/10.1145/3133956.3133982>.
- [16] S. Lundberg, S.-I. Lee, "A unified approach to interpreting model predictions," in NeurIPS, 2017. arXiv: 1705.07874.
- [17] Y. Zhang et al., "Graph neural networks in healthcare: a survey," IEEE J. Biomed. Health Inform., 2021. DOI: 10.1109/JBHI.2021.3050379.



- [18] E. J. Topol, “High-performance medicine: the convergence of human and artificial intelligence,” Nat. Med., 2019. DOI: 10.1038/s41591-0190447-x.
- [19] Regulatory perspectives on AI in healthcare (FDA/EMA guidance summaries), 2020. Example: <https://www.fda.gov/medical-devices/softwaremedical-device-samd>.
- [20] D. Saria, “Towards clinically actionable AI: evaluation, deployment and trust,” J. Clin. Transl. Sci., 2021. DOI: 10.1017/cts.2021.9.
- [21] Nutritionix resources and dataset access, 2020–2023. URL: <https://www.nutritionix.com/>.
- [22] A. Esteva et al., “A guide to deep learning in healthcare,” Nat Med, 2019. DOI: 10.1038/s41591-019-0447-x.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)