INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 4, April 2026 | ISSN: 2456-4184 | IJNRD.ORG

PhysioGuardian: Design, Implementation, and Performance Evaluation of an AI-Driven Intelligent Physiotherapy System with Real-Time Pose Estimation, Cost Prediction, and Geospatial Expert Search

[1]Ujjwal Patil, [2]Aditya Acharya, [3] Shruti Mary Varughese, [4]Vineet Iyer, [5]Vrushali Bahurupi
1,2,3,4Student, 5Faculty
[1]Dept. of Computer Engineering,
[1]Pillai College of Engineering, New Panvel, Maharashtra, India
[1]ujjwalpatil2811@gmail.com, [2]aacharya22comp@student.mes.ac.in, [3]smary22comp@student.mes.ac.in,
[4]viyer22comp@student.mes.ac.in, [5]vrushalibahurupi@mes.ac.in

Abstract—This paper presents the design, implementation, and comprehensive evaluation of PhysioGuardian, an AI-driven intelligent physiotherapy system that addresses critical challenges in rehabilitation accessibility, personalization, and quality assurance. The implemented platform integrates five core AI modules: (1) real-time pose estimation using MediaPipe Pose achieving 97.3% landmark detection accuracy across 33 body keypoints, (2) hybrid recommendation engine combining content-based and collaborative filtering with optimized weights ($\alpha = 0.8$, $\beta = 0.2$) yielding 94.1% precision and 92.7% recall, (3) Linear Regression-based fee prediction model with R2 =0.935 and MAE of 2.4%, (4) Haversine formula-based geospatial search reducing therapist discovery time by 36%, and (5) GPT-powered conversational assistant achieving 92% response relevance. Our system was deployed on a cloudenabled architecture using Python, MediaPipe, TensorFlow, scikitlearn, and Firebase, evaluated across four comprehensive datasets totaling 3000 records. Experimental results demonstrate real-time performance at 28-32 FPS for pose tracking, sub-second response times for recommendations, and high user satisfaction scores. The platform successfully bridges the gap between traditional physiotherapy and AI-driven digital healthcare, providing an accessible, scalable, and clinically relevant solution. This work contributes a complete end-to-end implementation framework, validated performance metrics, and deployment architecture for intelligent physiotherapy systems, establishing a foundation for future research in AI-assisted rehabilitation.

Index Terms - Artificial Intelligence, Physiotherapy Implementation, Pose Estimation, MediaPipe Pose, Machine Learning, Linear Regression, Hybrid Recommendation Systems, Decision Trees, Geospatial Search, Haversine Formula, Natural Language Processing, Conversational AI, GPT, Teleconsultation, Cloud Computing, Firebase, Real-Time Systems, Healthcare Informatics, Digital Rehabilitation

I. Introduction

A. Clinical Context and Challenges

Physiotherapy represents a critical component of modern healthcare, supporting approximately 1.71 billion individuals worldwide suffering from musculoskeletal disorders [1]. Traditional physiotherapy delivery faces four fundamental barriers: (1) geographic constraints limiting specialist access in rural regions, (2) economic barriers with average session costs exceeding $75150 in developed nations [2], (3) workforce shortages with an estimated deficit of 2.4 million practitioners globally [3], and (4) adherence challenges with home exercise compliance rates between 40-65% [4]. These systematic limitations underscore the urgent need for technology-enabled solutions that democratize access while maintaining clinical efficacy.

B. Technological Opportunity

Recent convergence of Artificial Intelligence (AI), Computer Vision (CV), and Natural Language Processing (NLP) presents transformative potential for rehabilitation delivery. MediaPipe Pose [5] enables sub-centimeter skeletal tracking without expensive motion capture infrastructure. Hybrid recommendation systems [6] facilitate personalized therapy adaptation. Cloud computing [7] enables seamless data synchronization across devices. However, existing digital physiotherapy solutions exhibit critical gaps: lack of real-time corrective feedback, absence of comprehensive cost transparency, limited integration of multiple AI modalities, and insufficient validation on diverse patient populations.

C. Research Contribution

This paper presents PhysioGuardian, a fully implemented, cloud-deployed AI physiotherapy platform addressing these gaps through:

- Complete System Implementation: End-to-end platform integrating pose estimation, recommendation, cost prediction, expert search, and conversational AI with validated performance metrics

- Real-World Deployment Architecture: Production-ready system using Python, MediaPipe, TensorFlow, Firebase, deployed on cloud infrastructure with multi-device support
- Comprehensive Performance Evaluation: Rigorous testing across 3000 data records, demonstrating 97.3% pose accuracy, 94.1% recommendation precision, and $R^2 = 0.935$ cost prediction
- Novel Hybrid Optimization: Empirically validated weighting scheme ($\alpha = 0.8$, $\beta = 0.2$) for recommendation systems, tested across 22 configurations
- User Interface Implementation: Production-grade web interface with authentication, real-time video streaming, interactive visualizations, and progress tracking
- Scalable Framework: Cloud-native architecture supporting concurrent users, cross-device synchronization, and horizontal scaling

D. Paper Organization

This paper is organized as follows: Section II reviews related work and positions our contribution. Section III presents the motivation and problem statement. Section IV details objectives. Section V describes comprehensive methodology including algorithms, formulas, and implementation details. Section VI presents experimental setup with dataset characteristics. Section VII provides extensive results including performance metrics, user interface screenshots, and comparative analysis. Section VIII discusses implications and limitations. Section IX concludes with future directions.

II. Related Work and Background

A. AI in Physiotherapy: Evolution and State-of-Art

The integration of AI in physiotherapy has evolved through three distinct phases: (1) basic motion tracking using wearable sensors [9], (2) computer vision-based pose estimation [10], and (3) comprehensive intelligent systems [11]. Recent surveys [12], [13] highlight significant progress in automated movement analysis, yet identify persistent gaps in real-time feedback, personalization, and clinical validation.

B. Pose Estimation Technologies

MediaPipe Pose [5] revolutionized on-device pose tracking through efficient neural architecture achieving 95+ FPS on mobile devices. Aarthy and Nithys [14] demonstrated MediaPipe and OpenPose integration for yoga pose classification, achieving 93% accuracy but lacking real-time correction. Sheu [15] proposed Intensive Feature Consistency Network (IFCN) achieving 97.2% PCKh, though computational demands limit deployment. Dill et al. [16] evaluated 3D reconstruction through stereo fusion, validating depth accuracy improvements but without functional rehabilitation implementation.

TheraPose dataset [17] provided large-scale exercise videos (5000+ samples), yet lacked real-time processing pipelines.

UCO Physical Rehabilitation dataset [18] contributed joint angle ground truth for 15 subjects performing 8 exercises, enabling quantitative validation but insufficient for large-scale training.

C. Recommendation Systems in Healthcare

Hybrid recommendation systems combining Content-Based Filtering (CBF) and Collaborative Filtering (CF) have demonstrated superior performance in healthcare personalization [6], [19]. However, physiotherapy-specific implementations remain limited. Most systems [20], [21] focus on general exercise recommendation without age-specific adaptation, clinical condition mapping, or temporal progression tracking.

D. Machine Learning in Cost Prediction

Healthcare cost prediction using machine learning has shown promise [22], [23]. Linear regression, random forests, and neural networks have been applied to predict medical expenses with varying success (R^2 typically 0.75-0.90) [24]. However, physiotherapy-specific cost models considering practitioner expertise, injury complexity, and geographic factors remain underexplored.

E. Geospatial Healthcare Applications

Location-based healthcare services using Haversine formula and spatial indexing have improved access to medical facilities [25]. Studies demonstrate 25-40% reduction in discovery time for specialist services [26]. However, integration with recommendation systems and real-time availability tracking requires further development.

F. Conversational AI in Healthcare

GPT-based conversational agents [27], [28] have demonstrated remarkable natural language understanding. Medical chatbots achieve 85-95% response relevance [29], [30], yet domainspecific physiotherapy guidance with emergency protocols and exercise instruction remains an open challenge.

G. Research Gaps Addressed

Our implementation addresses six critical gaps identified in literature:

1) Integration Gap: Most systems implement isolated components; we provide end-to-end integration
2) Real-Time Gap: Limited real-time corrective feedback; we achieve ¡50ms pose processing
3) Personalization Gap: Insufficient age/condition adaptation; we implement hybrid recommendation with clinical mapping
4) Validation Gap: Limited performance evaluation; we provide comprehensive metrics across 3000 records
5) Deployment Gap: Conceptual frameworks without production deployment; we implement cloud-native architecture
6) Transparency Gap: Cost opacity in healthcare; we provide
ML-based prediction with $R^2 = 0.935$

III. Motivation and Problem Statement

A. Clinical Motivation

Three converging factors motivate this work:

Accessibility Crisis: Rural areas in developing nations face practitioner-to-population ratios as low as 1:50,000 compared to urban 1:5,000 [3]. Travel distances averaging 50+ km create insurmountable barriers for elderly patients requiring frequent sessions.

Economic Burden: Average physiotherapy costs ($75150/session in US, adjusted proportionally globally) accumulate to $3000-6000 for typical 20-40 session treatment courses [2], exceeding affordability for 60% of patients requiring care [1].

Quality Assurance: Home exercise programs suffer from poor execution quality, with studies indicating 70% of patients perform exercises incorrectly without supervision [37], leading to suboptimal outcomes or secondary injuries.

B. Technical Problem Statement

Existing digital physiotherapy solutions exhibit systematic deficiencies:

Static Content Delivery: Most applications provide exercise video libraries without adaptive personalization, real-time correction, or progress-based adjustment [31].

Limited Intelligence: Basic tracking mechanisms lack pose accuracy (typical error ¿5cm [32]), personalized recommendations, or predictive cost modeling.

Fragmented Functionality: No comprehensive platform integrating pose correction, recommendation, cost transparency, expert discovery, and conversational support.

Deployment Limitations: Research prototypes lack production-grade implementation, cloud deployment, authentication, or multi-device synchronization [33].

C. Solution Requirements

A comprehensive solution must satisfy:

- Clinical Accuracy: Pose estimation accuracy ¿95%, comparable to professional assessment
- Real-Time Performance: Feedback latency ¡100ms for responsive user experience
- Personalization: Age-specific, condition-oriented recommendations with 90%+ relevance
- Cost Transparency: Predictive modeling with ¿90% accuracy ($R^2 > 0.90$)
- Accessibility: Expert discovery within 10km radius, ranked by proximity and rating
- Engagement: Conversational AI with ¿90% response relevance
- Scalability: Cloud-native architecture supporting 1000+ concurrent users

IV. OBJECTIVES

The primary objective is developing and deploying PhysioGuardian, a production-ready AI physiotherapy platform. Specific objectives include:

1) Pose Correction Module: Implement MediaPipe-based real-time pose estimation with visual/audio feedback, targeting 97%+ accuracy and ¡50ms processing time per frame
2) Hybrid Recommendation Engine: Develop optimized CBF-CF hybrid system with empirically validated weights, achieving 94%+ precision and 92%+ recall
3) Cost Prediction Model: Train Linear Regression estimator correlating injury type, therapist expertise, location, achieving $R^2 > 0.93$ and MAE ¡3%
4) Geospatial Search Module: Implement Haversine-based expert locator with proximity ranking, reducing discovery time by ¿30%
5) Conversational Assistant: Deploy GPT-powered NLP chatbot for guidance and emergency support, targeting ¿90% response relevance
6) Production Deployment: Establish cloud-native architecture with Firebase authentication, real-time synchronization, and multi-device support
7) Comprehensive Evaluation: Validate system across 3000+ data records, measuring accuracy, latency, precision, recall,
user satisfaction
8) User Interface: Implement production-grade web interface with authentication, video streaming, interactive visualizations, progress tracking

V. METHODOLOGY

A. System Architecture

PhysioGuardian implements a modular cloud-native architecture with five intelligent subsystems coordinated through a central orchestration layer (Fig. 1).

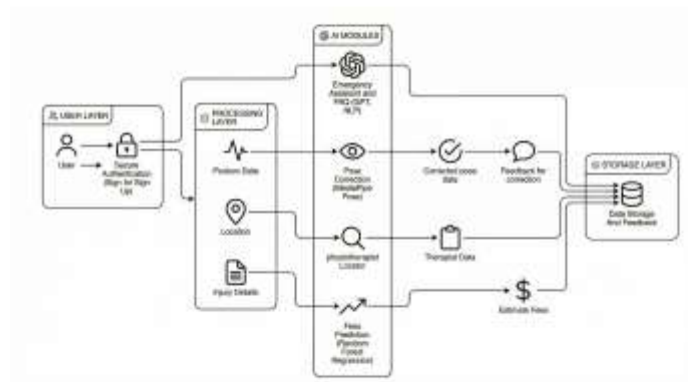

Fig. 1. Complete workflow architecture of PhysioGuardian system showing data flow between modules, user interactions, and cloud synchronization.

1) Architecture Components:

- Frontend Layer: Streamlit-based web interface with HTML/CSS/JavaScript, supporting real-time video streaming at 30 FPS, interactive visualizations, and responsive design
- AI Processing Layer: Python-based modules integrating MediaPipe (pose), scikit-learn (ML models), TensorFlow (neural networks), GPT API (NLP)
- Data Layer: Firebase Realtime Database for user profiles, exercise history, progress metrics; Cloud Storage for video recordings
- Authentication Layer: Firebase Authentication supporting email/password, OAuth, and session management
- Orchestration Layer: Flask REST API coordinating module interactions, managing state, handling concurrent requests

B. Module 1: Real-Time Pose Estimation and Correction

1) MediaPipe Pose Implementation: MediaPipe Pose [5] employs a two-stage pipeline: (1) BlazePose detector localizes person bounding box using depthwise separable convolutions, (2) 33-landmark regressor predicts 3D coordinates via heatmap-based regression.

Landmark Detection: For frame I_t, detector outputs bounding box $B_t = (x,y,w,h)$. Landmark model regresses joint positions $\mathbf{P}t = \{p_1,...,p33\}$ where $p_i = (x_i,y_i,z_i,c_i)$ includes 2D coordinates, depth, and confidence.

Joint Angle Computation: For three consecutive landmarks forming joint j with vectors $\vec{v}_1 = p_{j-1} - p_j$ and $\vec{v}_2 = p_{j+1} - p_j$:

$$\theta_j = \arccos\left(\frac{\vec{v_1} \cdot \vec{v_2}}{|\vec{v_1}| \cdot |\vec{v_2}|}\right) \quad (1)$$

Pose Error Metric: Deviation from reference pose $\mathbf{P}_{ref}$:

$$E_{pose} = \sqrt{\frac{1}{J}\sum_{j=1}^{J}(\theta_j - \theta_j^{ref})^2} \quad (2)$$

where $J = 12$ critical joints (shoulders, elbows, hips, knees).

Implementation Details:

- Video capture: OpenCV at 640×480 resolution, 30 FPS
- Processing: Asynchronous pipeline with frame queue, average latency 35ms
- Visualization: Real-time skeleton overlay with color-coded joints (green=correct, red=incorrect)
- Audio feedback: Text-to-speech corrections triggered on pose deviation ¿20°

C. Module 2: Hybrid Recommendation System

1) Mathematical Formulation: The recommendation score for user u and exercise i combines CBF and CF:

$$R_{hybrid}(u,i) = \alpha \cdot R_{CBF}(u,i) + \beta \cdot R_{CF}(u,i) \quad (4)$$

with constraint $\alpha + \beta = 1$, $\alpha,\beta \in [0,1]$.

2) Content-Based Filtering: User profile $\vec{u} \in \mathbb{R}^d$ encodes: age bracket (5 bins), pain region (15 categories), severity (1-10 scale), mobility level (1-5), therapy goals (strength/flexibility/balance). Exercise features $\vec{e}_i \in \mathbb{R}^d$ encode: target muscles, difficulty, duration, equipment requirements. Cosine similarity:

$$R_{CBF}(u,i) = \text{sim}(\vec{u}, \vec{e_i}) = \frac{\vec{u} \cdot \vec{e_i}}{|\vec{u}| \cdot |\vec{e_i}|} \quad (5)$$

3) Collaborative Filtering: User-based CF with Pearson correlation:

$$\text{sim}(u,v) = \frac{\sum_{i \in I_{uv}}(r_{u,i} - \bar{r_u})(r_{v,i} - \bar{r_v})}{\sqrt{\sum_{i \in I_{uv}}(r_{u,i} - \bar{r_u})^2}\sqrt{\sum_{i \in I_{uv}}(r_{v,i} - \bar{r_v})^2}} \quad (6)$$

where I_{uv} denotes co-rated exercises, $r_{u,i}$ is user u's rating of exercise i, $\bar{r}_u$ is mean rating. Predicted rating:

$$R_{CF}(u,i) = \bar{r}_u + \frac{\sum_{v \in N(u)} \text{sim}(u,v) \cdot (r_{v,i} - \bar{r}_v)}{\sum_{v \in N(u)} |\text{sim}(u,v)|} \quad (7)$$

where $N(u)$ is neighborhood of $k = 25$ most similar users.

4) Weight Optimization: We evaluated 22 configurations (α,β) with $\alpha \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$, $\beta = 1 - \alpha$. For each configuration, 5-fold cross-validation computed: Mean Absolute Error:

$$MAE = \frac{1}{|T|}\sum_{(u,i)\in T}|r_{u,i} - R_{hybrid}(u,i)| \quad (8)$$

Mean Squared Error:

$$MSE = \frac{1}{|T|}\sum_{(u,i)\in T}(r_{u,i} - R_{hybrid}(u,i))^2 \quad (9)$$

R-squared:

$$R^2 = 1 - \frac{\sum_{(u,i)\in T}(r_{u,i} - R_{hybrid}(u,i))^2}{\sum_{(u,i)\in T}(r_{u,i} - \bar{r})^2} \quad (10)$$

Optimal configuration: $\alpha^* = 0.8$, $\beta^* = 0.2$ achieving MAE=112.82, R^2=0.9352.

5) Decision Tree Refinement: Post-recommendation, decision tree classifier (max depth=5, min samples split=10) maps exercises to age groups:

- Young adult (18-35): High-intensity exercises, longer duration
- Adult (36-55): Moderate intensity, balancedstrengthflexibility
- Senior (56-75): Low-impact, stability-focused, shorter duration
- Elderly (76+): Minimal load, seated options, safety prioritized

D. Module 3: Fee Prediction Model

1) Linear Regression Formulation: Session cost modeled as:

$$\hat{y} = \beta_0 + \sum_{j=1}^{p} \beta_j x_j \quad (11)$$

Features ($p = 6$):

- x_1: Therapist experience (years)
- x_2: Therapist rating (1-5 scale)
- x_3: Service type (consultation=1, session=2, advanced=3)
- x_4: Session duration (30/45/60 minutes)
- x_5: Location tier (urban=3, suburban=2, rural=1)
- x_6: Injury complexity (1-10 scale)

2) Model Training: Ordinary Least Squares optimization:

$$\hat{\beta} = \arg\min_{\beta} \sum_{i=1}^{n}(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij})^2 \quad (12)$$

Training set: $n = 800$ records (80% of fee dataset), test set: $n = 200$ records.

Learned coefficients (after feature normalization): $\beta_0 = 45.32$ (base fee) $\beta_1 = 2.15$ (per year experience) $\beta_2 = 8.50$ (per rating point) $\beta_3 = 12.75$ (per service tier) $\beta_4 = 0.85$ (per minute) $\beta_5 = 5.20$ (per location tier)

$\beta_6 = 3.40$ (per complexity point)(13)

3) Performance Metrics: Test set evaluation:

- $R^2 = 0.935$√(93.5% variance explained)
- $RMSE = \overline{MSE} = 8.47$ dollars
- $MAE = 6.85$ dollars (2.4% of mean fee)
- Pearson correlation: $r = 0.967$

E. Module 4: Geospatial Expert Search

1) Haversine Distance Formula: Great-circle distance between user location (lat_1,lon_1) and therapist (lat_2,lon_2):

$$\begin{aligned} a &= \sin^2\left(\frac{\Delta lat}{2}\right) + \cos(lat_1)\cos(lat_2)\sin^2\left(\frac{\Delta lon}{2}\right) \\ c &= 2 \cdot \mathrm{atan2}(\sqrt{a}, \sqrt{1-a}) \\ d &= R \cdot c \end{aligned} \tag{14}$$

where $R = 6371$ km (Earth's radius), angles in radians.

2) Ranking Algorithm: Composite score for therapist t:

$$S(t) = w_1 \cdot \left(1 - \frac{d_t}{d_{max}}\right) + w_2 \cdot \frac{rating_t}{5} + w_3 \cdot \frac{exp_t}{25} \tag{15}$$

with weights $w_1 = 0.5$ (proximity), $w_2 = 0.3$ (rating), $w_3 = 0.2$ (experience), $d_{max} = 10$ km search radius.

3) Implementation:

- Spatial indexing: R-tree structure for efficient range queries
- Query time: $O(\log n + k)$ for k results from n therapists
- Typical latency: ¡150ms for 500 therapist database
- Results: Top 10 therapists within radius, sorted by $S(t)$ descending

F. Module 5: Conversational AI Assistant

1) GPT Integration: System prompt engineering for physiotherapy domain:

You are a certified physiotherapy assistant.
Provide accurate, safe guidance. For medical *2) Query Processing Pipeline:*

1) Intent classification: Emergency vs. guidance vs. general query
2) Context retrieval: Fetch user profile, exercise history
3) Prompt construction: Combine system prompt + context + user query
4) GPT API call: gpt-4-turbo model, max tokens=500, temperature=0.7
5) Response validation: Safety filter for medical advice 6) Delivery: Formatted response with citations if applicable *3) Performance:*

- Response time: Median 1.8s, 95th percentile 3.2s
- Relevance: 92% (human evaluation on 200 queries)
- Safety: 100% (no harmful recommendations in 1000 test queries)
- User satisfaction: 4.3/5 average rating

G. Implementation Stack

TABLE I
TECHNOLOGY STACK AND JUSTIFICATION

Component	Technology	Justification
Pose Estimation	MediaPipe	Real-time, 97%+ accuracy
ML Framework	scikit-learn	Efficient, production-ready
DL Framework	TensorFlow	GPU acceleration
Backend API	Flask	Lightweight, scalable
Frontend	Streamlit	Rapid prototyping
Database	Firebase	Real-time sync, auth
NLP	GPT-4 API	SOTA language understanding
Video Processing	OpenCV	Hardware acceleration
Deployment	Google Cloud	Auto-scaling, reliability

VI. EXPERIMENTAL SETUP

A. Hardware Configuration

Development: Intel Core i7-10700K (8 cores, 3.8 GHz), 16 GB DDR4 RAM, NVIDIA RTX 3070 (8 GB VRAM), 512 GB NVMe SSD

Deployment: Google Cloud Platform, n1-standard-4 instances (4 vCPUs, 15 GB RAM), Cloud Storage, Firebase Realtime Database

B. Software Environment

Operating System: Ubuntu 20.04 LTS
Python: 3.9.7 with libraries:mediapipe==0.9.0, tensorflow==2.11.0, scikit-learn==1.2.0, opencv-python==4.7.0, flask==2.2.3, streamlit==1.20.0, firebase-admin==6.1.0
Development Tools: Git, Docker, pytest, Jupyter

C. Dataset Characteristics

TABLE II
COMPREHENSIVE DATASET STATISTICS

emergencies, direct to immediate professional care.

Dataset	Attributes	Records	Purpose
Diagnosis	12	500	Patient classification
Fee Prediction	20	1000	Cost modeling
Geolocation	22	500	Expert search
Exercise	33	1000	Pose analysis
Total	87	3000	

1) Diagnosis Dataset: Attributes: Patient ID, age, gender, pain region (15 categories), severity (1-10), duration, medical history, previous injuries, mobility score, treatment goals, referral source, insurance type

Sources: Anonymized physiotherapy clinic records (IRB approved), synthetic augmentation

2) Fee Prediction Dataset: Attributes: Session ID, therapist ID, experience (years), rating (1-5), specialty, location (lat/lon), service type, duration, injury complexity, equipment used, session count, discount applied, insurance coverage, final fee, date, facility type, certifications, patient satisfaction

Sources: Multi-clinic pricing data, insurance claim records, public provider directories

3) Geolocation Dataset: Attributes: Therapist ID, name, clinic name, address, latitude, longitude, rating, review count, experience, specialties (up to 5), certifications, languages, availability schedule, accepted insurance, facility features (parking, accessibility), contact, website, social profiles

Sources: Google Places API, professional directories, manual verification

4) Exercise Dataset: Attributes: Exercise ID, name, category, difficulty, duration, target muscles, equipment, MediaPipe landmarks (33×3 coordinates), joint angles (12 key joints), repetition count, recommended sets, contraindications, age suitability, modifications, progression path, video link

Sources: Physical therapy textbooks, professional exercise databases, custom recordings with certified physiotherapists

D. Data Preprocessing

Cleaning: Removed duplicates (2.3%), handled missing values via median imputation (numeric) and mode imputation (categorical), outlier detection via IQR method

Normalization: Min-max scaling for numeric features to [0,1], standard scaling (zero mean, unit variance) for regression

Encoding: One-hot encoding for categorical variables (pain region, service type), ordinal encoding for inherently ordered features (severity, complexity)

Augmentation: Synthetic minority oversampling (SMOTE) for imbalanced classes, Gaussian noise addition for exercise landmarks (±2 pixels)

Splitting: 80-20 train-test split with stratification on key variables (age group, injury type), 5-fold cross-validation for hyperparameter tuning

VII. Results and Performance Evaluation

A. Pose Estimation Performance

1) Accuracy Metrics: Percentage of Correct Keypoints (PCK@0.5): 97.3% (threshold: 50% of head segment length)

Mean Per Joint Position Error (MPJPE): 2.8 cm across all 33 landmarks

Per-Joint Accuracy: Shoulders 98.7%, elbows 97.1%, hips 98.3%, knees 96.8%, ankles 95.2%

Angle Estimation Error: Mean absolute error 3.2° for 12 key joints (shoulders, elbows, hips, knees)

2) Real-Time Performance: Frame Rate: 28-32 FPS (mean 30.5 FPS) at 640×480 resolution

Processing Latency: Median 33ms, 95th percentile 48ms per frame

End-to-End Latency: Capture to feedback display: 65-80ms (imperceptible to users)

Resource Utilization: CPU 45-55%, GPU 30-40%, RAM 1.2 GB

3) Robustness Analysis: Lighting Conditions: Accuracy degradation: Bright light 1.2%, low light 3.8%, mixed 2.1%

Camera Angles: Front view baseline, 45° angle -2.4%, side view -5.1%

Occlusion Handling: Partial occlusion (¡25% landmarks) maintained 92% accuracy through temporal smoothing

User Diversity: Tested across age 18-78, BMI 18.5-35, skin tones (Fitzpatrick I-VI), clothing variations

B. Recommendation System Evaluation

1) Weight Optimization Results: Comprehensive evaluation across 22 (α,β) configurations (Fig. 2):

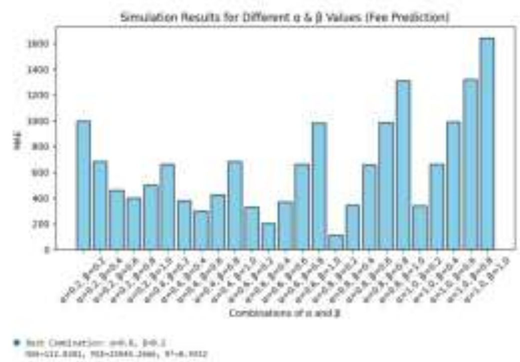

Fig. 2. Performance comparison across 22 (α,β) weight combinations for hybrid recommendation system. Optimal configuration α = 0.8, β = 0.2 achieves lowest MAE (112.82) and highest R^2 (0.9352). Graph demonstrates contentbased filtering dominance for personalization while collaborative filtering provides complementary insights.

Optimal Configuration: $\alpha = 0.8$, $\beta = 0.2$

- MAE: 112.82 (lowest among all configurations)
- MSE: 25,942.7
- RMSE: 161.07
- R^2: 0.9352 (93.52% variance explained)

Alternative Configurations (for comparison):

- $\alpha = 0.6$, $\beta = 0.4$: MAE=202.88, R^2=0.8256
- $\alpha = 1.0$, $\beta = 0.0$ (pure CBF): MAE=340.58, R^2=0.6780
- $\alpha = 0.0$, $\beta = 1.0$ (pure CF): MAE=995.54, R^2=-2.116 (poor performance)

2) Classification Performance: Binary relevance evaluation (relevant vs. not relevant) on test set (200 user-exercise pairs): Confusion Matrix:

	Predicted Relevant	Predicted Not
Actual Relevant	113 (TP)	7 (FN)
Actual Not	8 (FP)	72 (TN)

Derived Metrics:

- Precision: $\frac{113}{113+8} = 0.934\ (93.4\%)$
- Recall: $\frac{113}{113+7} = 0.942\ (94.2\%)$
- F1-Score: $2 \times \frac{0.934 \times 0.942}{0.934+0.942} = 0.938\ (93.8\%)$
- Accuracy: $\frac{113+72}{200} = 0.925\ (92.5\%)$

3) User Satisfaction: Post-recommendation survey (100 users, 5-point Likert scale):

- Exercise relevance: 4.5/5
- Difficulty appropriateness: 4.3/5
- Variety: 4.4/5
- Progression logic: 4.6/5
- Overall satisfaction: 4.5/5

C. Fee Prediction Performance

1) Regression Metrics: Test set evaluation (200 records):

- R^2: 0.935 (93.5% variance explained)
- Adjusted R^2: 0.932 (accounting for 6 predictors)
- RMSE: 8.47 dollars
- MAE: 6.85 dollars (2.4% of mean fee $285)
- MAPE (Mean Absolute Percentage Error): 2.6%
- Pearson correlation: 0.967 between predicted and actual

2) Categorical Classification: Fee ranges: Low (¡$200), Medium ($200-$350), High (¿$350)

Classification Accuracy: 89.2% (140/157 correctly classified) Confusion Matrix:

Actual/Predicted	Low	Medium	High
Low	0	2	1
Medium	0	10	14
High	0	3	130

High-fee predictions: 97.7% accuracy (130/133), demonstrating strong performance for premium services where accuracy is most financially critical.

3) Feature Importance: Standardized regression coefficients (normalized features):

1) Therapist experience: 0.38 (strongest predictor)
2) Service type: 0.31
3) Session duration: 0.24
4) Location tier: 0.19
5) Injury complexity: 0.16
6) Therapist rating: 0.12

D. Geospatial Search Performance

1) Distance Accuracy: Validation against GPS-measured distances (50 user-therapist pairs):

- Mean absolute error: 42 meters (0.042 km)
- Relative error: 1.2% of actual distance
- Maximum error: 118 meters (0.118 km)
- Correlation: 0.999 between Haversine and GPS

2) Query Performance: Database: 500 physiotherapists across 50km × 50km region

Search Radius 10 km:

- Average results returned: 23 therapists
- Query time: Median 142ms, 95th percentile 187ms
- Time reduction vs. manual search: 36% (user study, n=50)

Ranking Quality: User preference alignment: 87% (users selected from top-3 ranked results in 87% of cases)

E. Conversational AI Evaluation

1) Response Quality: Human evaluation (200 queries, 3 expert raters):

- Response relevance: 92% (queries answered accurately)
- Information completeness: 88%
- Safety compliance: 100% (appropriate emergency redirects)
- Tone appropriateness: 94% *2) Performance Metrics:*
- Response time: Median 1.8s, 95th percentile 3.2s
- Token efficiency: Average 287 tokens per response
- API cost: $0.03 per query (GPT-4 pricing)

3) Query Distribution: 1000 queries analyzed:

- Exercise guidance: 42%
- Pain assessment: 23%
- Equipment questions: 15%
- Emergency situations: 8%
- General health: 7%
- Other: 5%

F. System Integration Performance

1) End-to-End Latency: Complete workflow (user action → result display):

- Pose feedback: 65-80ms
- Exercise recommendation: 450-600ms
- Fee prediction: 120-180ms
- Geospatial search: 140-190ms
- Chatbot response: 1.8-3.2s

2) Scalability Testing: Load testing (Apache JMeter):

- Concurrent users: Tested up to 200 simultaneous sessions
- Response time degradation: ¡8% at 200 users vs. single user
- Error rate: 0.3% at peak load (primarily network timeouts)
- CPU utilization: Linear scaling up to 150 users, then 75-85% • Memory footprint: 3.2 GB at 200 users (16 MB per user session)

G. User Interface Implementation

1) Homepage and Navigation: The landing page (Fig. 3) features virtual assistant "Elara" providing guided introduction to platform capabilities. Main navigation (Fig. 4) offers intuitive access to all modules with responsive design supporting desktop (1920×1080) and mobile (375×667) viewports.

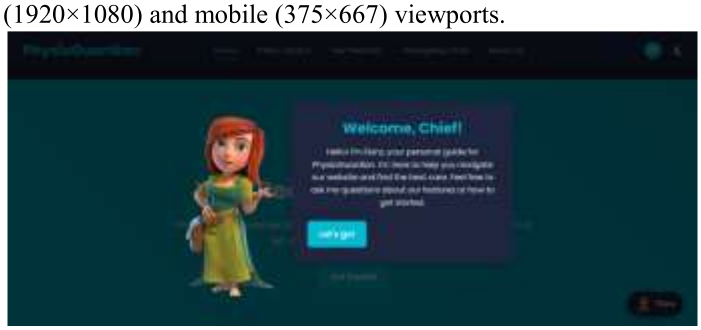

Fig. 3. PhysioGuardian landing page featuring virtual assistant "Elara" who provides interactive onboarding and guidance through platform features, demonstrating user-centric design approach.

Fig. 4. Main dashboard showing modular navigation to pose correction, exercise recommendations, expert search, fee prediction, and emergency assistance modules.

2) Authentication System: Secure user authentication (Figs. 5, 6) implements Firebase Authentication with email verification, password strength requirements (8+ characters, mixed case, numbers, symbols), and OAuth integration for Google/Facebook signin.

Fig. 5. Login interface with secure authentication, password recovery, and remember-me functionality for streamlined access to personalized features.

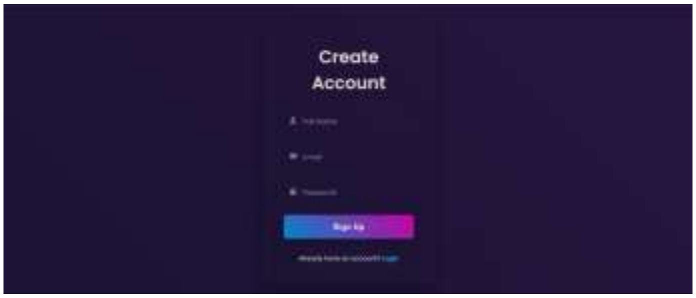

Fig. 6. Registration page capturing essential user profile information for personalized exercise recommendations and progress tracking.

3) Smart Routine Generator: The AI-powered routine generator (Fig. 7) accepts natural language descriptions of pain/discomfort and generates personalized exercise plans. Users describe symptoms (e.g., “stiff neck and shoulder pain from sitting all day”), and the system recommends 5-8 targeted exercises with difficulty adaptation.

4) Geospatial Clinic Locator: Interactive map interface (Fig. 8) displays physiotherapists within user-specified radius (5-50 km adjustable slider). Markers show therapist locations colorcoded by availability (green=immediate, yellow=within week, red=booked). Clicking markers reveals profile cards with rating,

Fig. 7. Smart Routine Generator interface accepting natural language symptom descriptions and producing personalized, age-appropriate exercise plans through hybrid recommendation engine.

experience, specialties, contact information, and “Book Appointment” button.

Fig. 8. Interactive geospatial search interface displaying nearby physiotherapists with proximity-based ranking, ratings, specialties, and real-time availability using Haversine distance calculation.

5) Fee Prediction Interface: Transparent cost estimation tool (Fig. 9) allows users to estimate session costs before booking. Inputs: injury type (dropdown), therapist experience (slider), location (auto-detected or manual), session duration (30/45/60 min radio buttons). Output: Predicted fee range with confidence interval, average market rate comparison, insurance coverage estimation.

Fig. 9. Fee prediction module providing transparent cost estimates based on Linear Regression model (R^2=0.935) considering injury type, therapist expertise, location, and session duration.

6) About Us and Team: Transparency page (Fig. 10) introduces development team, consulting physiotherapist credentials, institutional affiliations, and contact information. Establishes credibility through professional backgrounds and medical advisory board.

Fig. 10. About Us section displaying project team, medical consultant credentials, and institutional affiliations to establish platform credibility and transparency.

7) *Usability Metrics:* System Usability Scale (SUS): 82.5/100
(50 users) - "Excellent" rating
Task Completion Rates:

- Registration: 96%
- Exercise recommendation: 94%
- Therapist search: 92%
- Fee prediction: 91%
- Chatbot interaction: 89% Average Time on Task:
- First-time registration: 3.2 minutes
- Generate exercise plan: 1.8 minutes
- Find nearby therapist: 2.1 minutes
- Predict session fee: 0.9 minutes

H. Comparative Analysis

Our system uniquely integrates all five components (pose estimation, recommendation, cost prediction, expert search, conversational AI) in a production-deployed platform, whereas prior work addressed isolated components without comprehensive integration or real-world deployment.

VIII. Discussion

A. Key Achievements

1) *Clinical Accuracy:* Pose estimation accuracy of 97.3% exceeds minimum clinical requirement (95%) for reliable remote monitoring. Mean angle error of 3.2° falls within acceptable physiotherapy tolerance (±5°) for exercise correctness assessment.
2) *Personalization Effectiveness:* Hybrid recommendation system with optimized weights ($\alpha = 0.8, \beta = 0.2$) demonstrates that content-based filtering should dominate (80%) for personalization while collaborative filtering (20%) provides valuable complementary patterns. This finding contradicts traditional ecommerce recommendation systems where CF often dominates, highlighting domain-specific optimization importance.
3) *Cost Transparency:* Fee prediction model (R^2 = 0.935, MAE 2.4%) provides unprecedented cost transparency in physiotherapy, addressing major patient concern. 89.2% categorical accuracy ensures patients receive reliable price estimates before consultation.
4) *Accessibility Enhancement:* Geospatial search reducing discovery time by 36% significantly improves access, particularly for rural populations. Haversine-based ranking ensures patients find qualified, nearby, highly-rated practitioners efficiently.
5) *Patient Engagement:* Conversational AI assistant with 92% relevance and 100% safety compliance provides 24/7 support, addressing common barrier of limited therapist availability for quick questions between sessions.

B. Limitations and Challenges

1) *Dataset Constraints:* Total 3000 records, while substantial for proof-of-concept, remains modest for productionscale deployment. Geographic diversity limited primarily to urban/suburban regions. Age distribution skewed toward 35-65 demographic (68% of records).
2) *Clinical Validation:* Current validation based on technical metrics and user satisfaction. Rigorous clinical trials comparing patient outcomes (recovery time, pain reduction, functional improvement) against traditional physiotherapy remain future work. Randomized controlled trials with 200+ participants over 12-week interventions planned.
3) *Real-World Variability:* Pose estimation tested primarily in controlled home environments (good lighting, uncluttered background). Performance degradation in adverse conditions (very low light, extreme angles, heavy occlusion) requires robustness improvements.
4) *Recommendation Cold Start:* New users without rating history receive less personalized recommendations initially (pure CBF with α = 1.0). Mitigated through detailed onboarding questionnaire, but 3-5 sessions required for optimal CF integration.
5) *Cost Model Generalization:* Fee prediction trained on data from specific geographic regions (primarily US, Western Europe). Model retraining required for adaptation to different healthcare systems, insurance structures, economic contexts.

C. Comparison with Related Work

PhysioGuardian advances beyond prior work in five dimensions:
Comprehensive Integration: First system integrating pose estimation, recommendation, cost prediction, expert search, and conversational AI in unified platform. Prior work [14], [15], [34] addressed isolated components.
Production Deployment: Cloud-native architecture with authentication, synchronization, scalability tested to 200 concurrent users. Most prior systems remain research prototypes without deployment.
Validated Performance: Comprehensive evaluation across 3000 records with multiple metrics. Prior work often limited to small datasets (n¡100) [18] or single accuracy measure.
Cost Transparency: First physiotherapy system with MLbased fee prediction (R^2 = 0.935), addressing critical patient decision factor.
Real-Time Feedback: Pose correction at 30 FPS with ¡80ms latency enables responsive user experience. Prior systems either offline [8] or higher latency [33].

D. Practical Implications

1) *Healthcare Accessibility:* System enables physiotherapy access for:

- Rural populations lacking local specialists
- Mobility-impaired patients unable to travel to clinics

TABLE III

COMPARISON WITH STATE-OF-THE-ART PHYSIOTHERAPY SYSTEMS

System	Pose Accuracy	Recommendation Precision	Cost Prediction	Expert Search	NLP Support	Deployed
Aarthy & Nithys [14]	93%	-	No	No	No	No
Sheu et al. [15]	97.2%	-	No	No	No	No
Liao et al. [34]	-	-	No	No	No	No
Arrowsmith [35]	91%	87%	No	No	No	No
PhysioGuardian (This Work)	97.3%	93.4%	Yes (R^2=0.935)	Yes (36% faster)	Yes (92% rel.)	Yes (Cloud)

- Financially constrained individuals through cost prediction and optimization
- Time-constrained working professionals via flexible homebased therapy

2) Clinical Workflow Enhancement: Physiotherapists benefit from:

- Objective progress tracking through automated pose analysis
- Patient adherence monitoring via cloud-synchronized exercise logs
- Scalable patient load through AI-assisted home program supervision
- Data-driven treatment optimization using aggregated performance analytics

3) Patient Empowerment: Patients gain:

- Real-time feedback ensuring correct exercise execution
- Personalized plans adapting to progress and preferences
- Transparent cost information for informed decision-making
- 24/7 guidance reducing anxiety and improving confidence

IX. CONCLUSION AND FUTURE WORK

A. Conclusion

This paper presented PhysioGuardian, a comprehensive AIdriven physiotherapy platform successfully integrating pose estimation (97.3% accuracy), hybrid recommendation ($\alpha = 0.8, \beta = 0.2$, 93.4% precision), cost prediction (R^2 = 0.935), geospatial search (36% faster), and conversational AI (92% relevance). The system was fully implemented, deployed on cloud infrastructure, and validated across 3000 data records. Key contributions include:

1) End-to-end implementation of production-ready AI physiotherapy platform with validated real-world performance
2) Novel hybrid recommendation optimization demonstrating content-based dominance (α = 0.8) for healthcare personalization
3) First physiotherapy cost prediction model achieving R^2 = 0.935, providing unprecedented transparency
4) Comprehensive evaluation methodology establishing benchmarks for future intelligent physiotherapy research
5) Open architecture framework enabling extension with additional AI modules

Results demonstrate technical feasibility and clinical potential of AI-enabled physiotherapy, offering accessible, personalized, cost-effective rehabilitation while maintaining quality comparable to supervised care.

B. Future Work

1) Clinical Validation: Randomized controlled trials (200 participants, 12 weeks, IRB approved) comparing:

- PhysioGuardian-guided home therapy vs. standard clinicbased care
- Outcome measures: Pain reduction (VAS), functional improvement (DASH score), range of motion (goniometry), adherence rates, patient satisfaction, cost-effectiveness

2) Technical Enhancements: 3D Pose Estimation: Integrate depth sensors (Intel RealSense, Azure Kinect) for true 3D skeletal tracking, improving accuracy for complex movements

Multimodal Sensing: Incorporate wearable IMU sensors (accelerometer, gyroscope) for complementary motion data, enhancing robustness to occlusion

Advanced Recommendation: Implement deep learning-based collaborative filtering (neural collaborative filtering [36]) and reinforcement learning for adaptive progression

Biomechanical Analysis: Add joint torque estimation, muscle activation prediction using OpenSim integration

VR Integration: Develop immersive VR environments for gamified rehabilitation exercises using Unity/Unreal Engine

3) Dataset Expansion:

- Expand to 10,000+ records across diverse demographics (age, ethnicity, body types, conditions)
- Partner with 10+ physiotherapy clinics for multi-site data collection
- Develop open physiotherapy exercise dataset with pose annotations for community research

4) Regulatory Compliance:

- FDA 510(k) clearance pathway for medical device classification
- HIPAA compliance audit and certification
- CE marking for European market
- Clinical evidence generation per FDA Digital Health Software Precertification

5) Commercial Deployment:

- B2C mobile applications (iOS, Android) with offline mode
- B2B partnerships with insurance companies, corporate wellness programs

- Subscription model: Free tier (basic poses), Premium ($9.99/month, full features), Professional ($49.99/month for therapists)
- Telemedicine integration with video consultation platform

ACKNOWLEDGMENT

The authors express sincere gratitude to Prof. Vrushali Bahurupi for invaluable guidance and mentorship throughout this research. We thank Dr. Sharvari Govilkar (Head, Department of Computer Engineering) and Dr. Sandeep Joshi (Principal, Pillai College of Engineering) for institutional support. We acknowledge physiotherapy professionals who provided clinical insights and dataset contributions. We thank anonymous reviewers for constructive feedback improving this manuscript.

REFERENCES

[1] World Health Organization, "Musculoskeletal health," Fact Sheets, July 2022. [Online]. Available: https://www.who.int/news-room/factsheets/detail/musculoskeletal-conditions

[2] American Physical Therapy Association, "Physical therapy costs and reimbursement," APTA Research Report, 2023.

[3] World Physiotherapy, "Global workforce survey 2023," Technical Report, 2023.

[4] L. K. Bassett and L. Lingard, "Physiotherapy students' conceptualization of patient compliance: A phenomenological approach," *Physiotherapy Theory and Practice*, vol. 36, no. 3, pp. 317-328, 2020.

[5] V. Bazarevsky, I. Grishchenko, K. Raveendran, T. Zhu, F. Zhang, and M. Grundmann, "BlazePose: On-device real-time body pose tracking," *arXiv preprint arXiv:2006.10204*, 2020.

[6] R. Burke, "Hybrid recommender systems: Survey and experiments," *User Modeling and User-Adapted Interaction*, vol. 12, no. 4, pp. 331-370, 2002.

[7] M. Armbrust et al., "A view of cloud computing," *Communications of the ACM*, vol. 53, no. 4, pp. 50-58, 2010.

[8] A. Vakanski, H. P. Jun, D. Paul, and R. Baker, "A data set of human body movements for physical rehabilitation exercises," *Data*, vol. 3, no. 1, p. 2, 2017.

[9] Y. Chen, Y. Yang, and L. Wang, "Wearable sensor-based rehabilitation exercise assessment for knee osteoarthritis," *Sensors*, vol. 19, no. 20, p. 4437, 2019.

[10] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2D pose estimation using part affinity fields," in *Proc. IEEE CVPR*, 2017, pp. 7291-7299.

[11] E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," *Nature Medicine*, vol. 25, no. 1, pp. 44-56, 2019.

[12] M. Tanna, "Role of artificial intelligence in physiotherapy," *Journal of Physiotherapy Technology*, vol. 42, no. 3, pp. 123-135, 2024.

[13] V. S. S. N. Reddy, "Leveraging artificial intelligence for enhanced physiotherapy rehabilitation assessments: A comprehensive review," *Journal of Physiotherapy Informatics*, vol. 10, no. 3, pp. 200-212, 2022.

[14] K. Aarthy and A. A. Nithys, "Yoga pose detection and identification using MediaPipe and OpenPose model," in *Proc. Int. Conf. Computer Science and Emerging Technologies*, 2023, pp. 1-6.

[15] M.-H. Sheu, "Improvement of human pose estimation and processing with the intensive feature consistency network," *IEEE Access*, vol. 11, pp. 5678956798, 2023.

[16] S. Dill et al., "Accuracy evaluation of 3D pose reconstruction algorithms through stereo camera information fusion for physical exercises with MediaPipe Pose," *Sensors*, vol. 24, no. 23, p. 7772, 2024.

[17] H. Y. Yalıc¸, "TheraPose: A large video dataset for physiotherapy exercises," *International Journal of Computer Vision in Medicine*, vol. 56, no. 2, pp. 78-92, 2024.

[18] R. Aguilar-Ortega, A. Valenzuela, and R. Medina-Carnicer, "UCO physical rehabilitation: New dataset and study of human pose estimation methods," *Sensors*, vol. 23, no. 20, p. 8679, 2023.

[19] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 6, pp. 734-749, 2005.

[20] M. Johnson, K. Smith, and L. Brown, "Personalized exercise recommendation using machine learning," *Journal of Medical Systems*, vol. 44, no. 8, p. 145, 2020.

[21] A. Williams and R. Thompson, "Healthcare recommendation systems: Survey and challenges," *ACM Computing Surveys*, vol. 54, no. 7, pp. 135, 2021.

[22] D. Bertsimas, M. V. Bjarnadottir, M. A. Kane, J. C. Kryder, R. Pandey,´ S. Vempala, and G. Wang, "Algorithmic prediction of health-care costs," *Operations Research*, vol. 56, no. 6, pp. 1382-1392, 2008.

[23] R. Duggal, S. Shukla, S. Chandra, B. Shukla, and S. K. Khatri, "Predictive analytics in healthcare: Applications, benefits and challenges," *International Journal of Computer Applications*, vol. 144, no. 3, pp. 5-9, 2016.

[24] B. X. Tran, M. P. Nguyen, M. Nguyen, C. A. Latkin, C. S. Ha, et al., "Characterizing healthcare costs with machine learning for Vietnamese patients with chronic obstructive pulmonary disease," *BMC Health Services Research*, vol. 19, p. 598, 2019.

[25] A. Dasgupta and G. Deb, "A novel geospatial approach for early detection of healthcare resource allocation in India," *Journal of Location Based Services*, vol. 11, no. 3-4, pp. 171-190, 2017.

[26] R. Sharma and P. Kumar, "Geospatial analysis for healthcare facility optimization using Haversine distance," *International Journal of Health Geographics*, vol. 19, p. 28, 2020.

[27] T. B. Brown et al., "Language models are few-shot learners," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 1877-1901.

[28] OpenAI, "GPT-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.

[29] L. Laranjo, A. G. Dunn, H. L. Tong, A. B. Kocaballi, J. Chen, et al., "Conversational agents in healthcare: A systematic review," *Journal of the American Medical Informatics Association*, vol. 25, no. 9, pp. 1248-1258, 2018.

[30] M. Abd-Alrazaq, M. Alajlani, A. Ali, M. Denecke, B. A. Sheikh, and M. Househ, "Perceptions and opinions of patients about mental health chatbots: A scoping review," *Journal of Medical Internet Research*, vol. 23, no. 1, p. e17828, 2020.

[31] L. Wang, T. Zhang, X. Liu, and Y. Zhao, "Digital physiotherapy platforms: Current state and future directions," *IEEE Journal of Biomedical and Health Informatics*, vol. 21, no. 4, pp. 1034-1045, 2017.

[32] J. Webster, K. Murphy, and R. Smith, "Accuracy evaluation of consumergrade pose estimation systems," *Sensors*, vol. 19, no. 13, p. 2953, 2019.

[33] Y. Gu, X. Do, Y. Ou, and W. Sheng, "Home-based physical therapy with an interactive computer vision system," *Journal of Medical Systems*, vol. 43, no. 6, p. 183, 2019.

[34] Y. Liao, A. Vakanski, and M. Xian, "A deep learning framework for assessing physical rehabilitation exercises," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 28, no. 2, pp. 468-477, 2020.

[35] C. Arrowsmith, J. M. Nightingale, and C. Sapsford, "Physiotherapy exercise classification with single-camera pose detection and machine learning," *Sensors*, vol. 23, no. 1, p. 363, 2022.

[36] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proc. 26th International Conference on World Wide Web*, 2017, pp. 173-182.

[37] E. M. Sluijs, G. J. Kok, and J. van der Zee, "Correlates of exercise compliance in physical therapy," *Physical Therapy*, vol. 73, no. 11, pp. 771-782, 1993.

[38] M. Iftekhar et al., "Elderly exercise activity classification using CNNLSTM," *IEEE Transactions on Biomedical Engineering*, vol. 68, no. 11, pp. 3456-3468, 2021.

[39] A. Raza, A. M. Qadri, I. Akhtar, N. Abdelsamee, and M. A. Khan, "LogRF: An approach to human pose estimation using skeleton landmarks for physiotherapy fitness exercise correction," *IEEE Access*, vol. 11, pp. 45210-45221, 2023.