AMERICAN JOURNAL OF MANAGEMENT AND IOT MEDICAL COMPUTING
Peer Reviewed, Referred & Indexed Journal
E-ISSN: 3069-0110
Vol.5, No.2(2) (2026)
www.ajmimc.com

REALESTATE HOUSEPRICE PREDICTION

Mr. Prateek kumar kar
Student, Dept. of CSE,
GIFT Autonomous, Bhubaneswar

Mr. Debaprasad Mahala
Student, Dept. of CSE,
GIFT Autonomous, Bhubaneswar

Prof. Mohapatra Girashree Sahu
Project Guide, Dept. of CSE,
GIFT Autonomous, Bhubaneswar

Abstract— The rapid growth of AI-driven real estate analytics and increasing complexity of property valuation have created significant demand for intelligent house price prediction systems. Traditional valuation methods rely heavily on manual analysis and market expertise, which are time-consuming and inconsistent. This paper presents the design and implementation of an AI-Based Realestate Houseprice Prediction System that leverages Machine Learning, Data Analysis, Regression Models, and predictive analytics to estimate property prices accurately. The system analyses features such as location, area, number of bedrooms, amenities, market trends, and property images to generate accurate predictions. A web-based dashboard allows users to input property details and instantly receive estimated house prices. Experimental results demonstrate improved prediction accuracy and efficient performance for real-world real estate applications.

Keywords— Multi-Agent AI, Retrieval-Augmented Generation, Speech Recognition, Computer Vision, Data Analysis, Deep Learning, Gradio, Real Estate Assistant, House Price Prediction, TensorFlow.

I. INTRODUCTION

A. Background

In the modern digital era, Artificial Intelligence has become one of the most rapidly growing technologies across sectors including education, banking, transportation, and real estate. AI-based real estate systems are enabling faster price prediction, property prediction, virtual assistance, and continuous health monitoring by combining Machine Learning, Deep Learning, Natural Language Processing (Data Analysis), and Computer Vision [1].

The increasing penetration of smartphones and internet connectivity has made it feasible to deploy intelligent real estate services accessible from any location. Despite major advances in property science, significant populations — particularly in rural and remote areas — continue to face barriers including shortage of agents, high consultation costs, long waiting times, and lack of real estate awareness. Delayed property guidance frequently results in severe health complications that early intervention could have avoided [2].

Traditional real estate systems depend primarily on in-person interaction, and existing digital health applications mostly support only text-based communication without leveraging data interaction, property image analysis, or intelligent house price prediction. To overcome these limitations, Realestate Houseprice Prediction develops an integrated real estate assistant combining Speech Recognition, Data Analysis, Computer Vision, and Multi-Agent AI to provide preliminary real estate guidance through data, text, and property image input.

B. Problem Statement

Real Estate accessibility continues to be a major challenge in many regions, especially in rural and underdeveloped areas. Many people do not receive timely property support due to shortage of agents, high consultation costs, long waiting times, and lack of nearby real estate facilities. Existing real estate applications provide only basic text-based support and do not effectively utilize modern AI technologies such as data interaction, property image analysis, and intelligent house price prediction [3]. Most systems lack the capability to understand property features through natural conversation or analyse visible property conditions through property images. There is a clear need for an intelligent system providing instant preliminary property guidance, supporting data, text, and property image-based communication, improving accessibility for rural and remote users while enabling early house price detection.

C. Motivation

The motivation behind this project stems from the increasing need for accessible, intelligent, and efficient real estate solutions. The lack of real estate facilities in rural areas, the shortage of qualified agents, and the importance of early house price detection create a strong imperative for AI-powered real estate assistants available anytime and anywhere. Advances in AI technologies — specifically Machine Learning, Data Analysis, Speech Recognition, and Computer Vision — have made it feasible to create systems capable of understanding human speech, analysing property images, and generating intelligent real estate responses. The growing availability of open-source frameworks and widespread smartphone use further inspire the creation of digital real estate platforms [4].

D. Objectives

The primary objectives of the Realestate Houseprice Prediction platform are as follows:

- To design and develop an AI-based multimodal real estate assistant supporting data, text, and property image input.
- To implement Speech Recognition for data-based property feature analysis and Data Analysis for understanding user queries.
- To integrate Computer Vision and Deep Learning for property image-based house price detection.
- To deploy a Multi-Agent pipeline for property feature analysis, house price prediction, and recommendation generation.
- To implement Retrieval-Augmented Generation (Machine Learning) for accurate, data-grounded real estate responses.
- To develop a Gradio-based interactive frontend with text-to-speech data output functionality.
- To demonstrate the practical application of AI technologies in preliminary real estate assistance.

E. Scope of the Project

The scope of this project includes the design and implementation of a full-stack AI-based real estate assistant capable of processing data,

Received: 20-04-2026 | Accepted: 29-04-2026 | Published: 04-06-2026 | www.ajmimc.com

AMERICAN JOURNAL OF MANAGEMENT AND IOT MEDICAL COMPUTING
Peer Reviewed, Referred & Indexed Journal
E-ISSN: 3069-0110
Vol.5, No.2(2) (2026)
www.ajmimc.com

text, and property image inputs and delivering structured real estate guidance through a multi-agent pipeline and interactive dashboard. The system focuses on preliminary property feature analysis, house price prediction, and real estate awareness, leveraging Machine Learning and LLM-based reasoning while keeping the system practical and deployable without heavy infrastructure requirements.

II. LITERATURE REVIEW

A. Existing Real Estate Analytics Systems

Traditional real estate analytics platforms such as EHR systems, property feature checkers, and telemedicine tools have become standard in real estate delivery. These platforms provide useful functionalities but rely heavily on manual input and human interpretation. Text-based chatbots and data assistants have gained popularity but typically focus on a single interaction modality and lack integrated multimodal capabilities combining data, property image analysis, and intelligent house price prediction simultaneously [5].

B. AI in Real Estate Analytics

Artificial Intelligence has emerged as a transformative technology in real estate analytics, enabling systems to analyse vast buyer data and detect house price patterns that are difficult to identify manually. Machine Learning improves house price prediction accuracy while Deep Learning models detect abnormalities in property property images — X-rays, MRI scans, and skin house price photographs — with clinically significant accuracy. Data Analysis enables real estate chatbots to understand and respond to property queries naturally. Research consistently demonstrates that AI-driven real estate systems improve insight quality and enhance early house price detection [6].

C. Retrieval-Augmented Generation (Machine Learning)

Retrieval-Augmented Generation, introduced by Lewis et al. (2020), enhances LLM outputs by grounding them in retrieved context from an external knowledge base. Unlike fine-tuned models, Machine Learning systems retrieve relevant documents at inference time, significantly reducing hallucination and improving factual accuracy. In real estate applications, Machine Learning systems retrieve house price-related data, price estimation suggestions, and property guidelines from curated knowledge bases, ensuring AI-generated real estate recommendations are anchored to verified property information rather than generic model outputs [7].

D. Multi-Agent AI Frameworks

Multi-agent systems involve multiple AI agents collaborating to perform distinct tasks. Research consistently demonstrates that agent specialization produces higher-quality outputs than single-agent approaches on complex tasks. In AI real estate systems, multi-agent frameworks allow separate agents to handle data processing, Data Analysis property feature analysis, property image analysis, house price prediction, and recommendation generation independently while communicating via a structured pipeline. This modular approach improves scalability, flexibility, and overall performance [8].

E. Research Gap and Proposed Contribution

From the reviewed literature, there is a clear gap in systems that combine Machine Learning-based property knowledge retrieval with multi-agent analytical pipelines, real-time multimodal interaction supporting data, text, and property image input simultaneously, data-enabled accessible real estate interaction for elderly and rural users, and a deployable full-stack architecture with minimal infrastructure requirements. To address this, this project proposes Realestate Houseprice Prediction — an integrated platform combining Speech Recognition, Data Analysis, Computer Vision, Deep Learning, Machine Learning-based knowledge retrieval, and a multi-agent pipeline deployed via a Gradio-based interactive frontend.

III. SYSTEM OVERVIEW

A. Proposed System

Realestate Houseprice Prediction is a web-based real estate assistant that allows users to record data property features, type real estate queries, or upload property images of visible property conditions. Users receive Data Analysis-based property feature analysis, AI-powered house price predictions, and actionable real estate suggestions with precautionary recommendations. They can also listen to data responses via text-to-speech output. The system provides a Gradio dashboard with a data recording interface, property image upload panel, speech-to-text output display, agent response panel, and data response playback.

B. System Architecture

The system follows a six-layer pipeline architecture: (1) Input Layer — data, text, and property image inputs captured through the Gradio frontend; (2) Processing Layer — Speech Recognition converts data to text; (3) Data Analysis Layer — property feature extraction and intent recognition applied to text input; (4) Vision Layer — Computer Vision and Regression Model models analyse uploaded property property images; (5) Agent Layer — multi-agent pipeline executes House Price Prediction and Recommendation Agents; (6) Presentation Layer — Gradio dashboard displays results with text-to-speech data response output.

C. Technology Stack

Layer	Technology	Purpose
Input	Gradio	Voice, text & property image UI
Voice	SpeechRecognition / Whisper	Speech-to-text conversion
Data Analysis	spaCy / NLTK	Property Feature extraction & Data Analysis
Computer Vision	OpenCV / Regression Model (TensorFlow)	Image house price detection
Prediction	Scikit-learn / PyTorch	House Price prediction models
Machine Learning	LangChain / ChromaDB	Property knowledge retrieval
Voice Output	gTTS / pyttsx3	Text-to-speech responses
Database	MongoDB / SQLite	Real Estate data storage

Table I: Technology Stack of Realestate Houseprice Prediction

D. Key Functional Modules

Received: 20-04-2026 | Accepted: 29-04-2026 | Published: 04-06-2026 | www.ajmimc.com

Module	Description	Technology
Voice Agent	Converts speech to text	SpeechRecognition, Whisper
Data Analysis Engine	Property Feature & intent extraction	spaCy, NLTK
Computer Vision	Property property image analysis	OpenCV, Regression Model, TensorFlow
House Price Prediction	Predicts house price from property features	Scikit-learn, PyTorch
Machine Learning Module	Retrieves property knowledge	ChromaDB, LangChain
Recommendation	Generates real estate guidance	LLM + Machine Learning
TTS Module	Converts response to data	gTTS, pyttsx3
Dashboard	Interactive user interface	Gradio, HTML/CSS

Table II: Key Functional Modules

E. Advantages of Proposed System

The Realestate Houseprice Prediction platform eliminates the need for physical property market visits for basic preliminary property feature checking through full multimodal pipeline automation. It provides data-grounded real estate suggestions via Machine Learning, reducing hallucination and increasing property response reliability. The modular multi-agent design ensures quality at each analytical stage. The platform is accessible to elderly and non-technical users through data interaction. The system is deployable on standard hardware without heavy GPU infrastructure requirements.

IV. METHODOLOGY

A. Overall Workflow

The system operates through structured steps. First, the user provides input through data recording, text entry, or property property image upload via the Gradio interface. Second, Speech Recognition converts audio to text. Third, the Data Analysis Engine preprocesses text, extracts property features, and identifies real estate intent. Fourth, Computer Vision analyses uploaded property property images using Regression Model models for house price pattern detection. Fifth, the House Price Prediction Agent synthesizes property feature and property image analysis data and generates house price predictions. Sixth, the Machine Learning module retrieves relevant property knowledge from ChromaDB. Seventh, the Recommendation Agent generates real estate suggestions and precautions. Results are displayed in the chat panel and delivered as data response via text-to-speech.

B. Data Collection and Handling

The system utilises real estate datasets for training and testing Machine Learning models. House Price property features datasets provide structured property feature-house price mappings for classification model training. Property property image datasets include labelled skin house price property images, wound classifications, and visible infection patterns used to train Regression Model models. Real Estate knowledge documents are preprocessed into text chunks, embedded using sentence-transformer models, and stored in ChromaDB for Machine Learning retrieval. The data handling process includes input validation, missing value handling, data normalisation, property image resizing and enhancement, and serialization of records for embedding and model training.

C. Machine Learning Implementation

The Machine Learning system forms the property knowledge retrieval backbone. During indexing, property knowledge documents are split into text chunks, embedded using the all-MiniLM-L6-v2 sentence-transformer model, and stored in ChromaDB with source metadata. During retrieval, the user query and extracted property features are embedded, cosine similarity search is performed against ChromaDB collections, and the top-k most semantically relevant property records are returned as structured context for the Recommendation Agent prompt.

Score Range	Similarity	Action
> 0.85	High relevance	Primary context
0.70 – 0.85	Medium relevance	Secondary context
< 0.70	Low relevance	Exclude from prompt

Table III: Machine Learning Retrieval Scoring Criteria

D. Multi-Agent Pipeline

The multi-agent pipeline consists of four specialised agents operating sequentially. The Voice Processing Agent receives audio input, applies noise reduction, and outputs transcribed text. The Data Analysis Agent preprocesses transcribed or typed text, applies named entity recognition for property feature identification, and outputs a structured property feature JSON. The Computer Vision Agent receives the uploaded property property image, applies preprocessing, and runs the trained Regression Model classification model to identify house price patterns with confidence scores. The House Price Prediction Agent receives the combined property feature JSON and property image classification output, queries the Machine Learning module for relevant property context, and generates house price predictions with real estate suggestions and precautionary recommendations.

E. Backend API and Frontend

The backend handles system logic, AI model communication, and house price prediction processing, implemented in Python using Flask or FastAPI. Backend APIs receive requests from the Gradio frontend, process real estate data, and communicate with AI modules. The Gradio frontend uses a modular layout with data recording interface, property image upload panel, speech-to-text display, agent response panel, and data response playback. Average token consumption per recommendation pipeline execution is approximately 1,800 tokens with response time under 2 seconds per agent call.

V. SYSTEM DESIGN

A. Use Case Diagram

The Use Case Diagram illustrates interaction between users and the Realestate Houseprice Prediction platform. Both General Users and Elderly Users can record data property features, enter text property features, upload property property images, view house price predictions, receive real estate recommendations, and play data responses. The Administrator role has additional privileges for managing property datasets, retraining AI models, and system configuration. The system connects to an external Property Knowledge Base for Machine Learning-based information retrieval during recommendation generation.

B. ER Diagram

The ER Diagram represents data entities and relationships. The User entity (user_id, name, age) initiates sessions in the Interaction entity (interaction_id, user_id FK, timestamp, input_type). Each interaction includes a VoiceInput record (data_id, audio_data, transcribed_text) and/or PropertyImage record (property image_id, property image_path). Property Feature records are extracted from data and text input. The House PricePrediction entity (prediction_id, house price_name, confidence_score) links to Recommendation records (recommendation_id, precautions, price estimation_suggestion). The PropertyKnowledge entity (knowledge_id, content, embedding) is referenced during Machine Learning retrieval.

C. Sequence Diagram

The Sequence Diagram illustrates step-by-step interaction during a user session. The user provides data or text input and optionally uploads a property property image via the Gradio UI. The Voice Processing Module converts audio to text and forwards it to the Data Analysis Engine. The Data Analysis Engine extracts property features and passes them to the House Price Prediction Agent. If an property image is uploaded, the Computer Vision Module processes it concurrently. The Prediction Agent queries the Machine Learning Module, which retrieves relevant records from ChromaDB. Results are displayed in the Gradio UI and converted to data output via text-to-speech.

D. System Architecture Diagram

The system architecture follows a layered pipeline structure. The browser-based Gradio UI captures data, text, and property image input. The Python Backend invokes the Voice Processing Module and Data Analysis Engine for text-based input, while the Computer Vision Module handles property image analysis concurrently. Processed outputs feed into the Multi-Agent Orchestrator managing sequential execution of the House Price Prediction and Recommendation Agents. The Machine Learning Module connects to ChromaDB for property knowledge retrieval. Responses flow back through the pipeline to produce price prediction output and data responses via text-to-speech.

VI. IMPLEMENTATION

A. Voice Processing Implementation

The data processing module is implemented using Python's SpeechRecognition library and OpenAI Whisper for high-accuracy transcription. It reads audio recorded via the Gradio microphone interface, applies noise reduction using the noisereduce library, and transcribes speech to text for downstream Data Analysis processing. Whisper is used for improved accuracy in noisy environments. The pipeline runs in under 1 second for audio clips up to 10 seconds in length.

B. Data Analysis Implementation

The Data Analysis module is implemented using spaCy and NLTK. It receives transcribed or typed text, applies tokenisation, stopword removal, and lemmatisation. Named Entity Recognition with a custom real estate entity model identifies property features, body parts, and property terms. Intent classification determines whether the user is describing property features, asking for information, or requesting property image analysis. Extracted property feature entities are serialised as a structured JSON for the House Price Prediction Agent. Processing time is under 200 milliseconds per query.

C. Computer Vision Implementation

The Computer Vision module is implemented using OpenCV for property image preprocessing and a Regression Model model trained using TensorFlow/Keras for property image classification. Input property images are resized to 224×224 pixels, normalised, and augmented during training. The Regression Model architecture follows a ResNet-inspired design trained on dermatological and wound classification datasets. The module outputs a house price class label with a confidence score. Image preprocessing and classification execute in under 1.5 seconds per property image.

D. Multi-Agent Pipeline Implementation

Agent	Input	Output
Voice Processing	Audio recording	Transcribed text
Data Analysis Agent	Text / transcription	Property Feature JSON
Computer Vision	Property property image	House Price class + score
House Price Prediction	Property Feature JSON + property image	Predictions + advice

Table IV: Multi-Agent Pipeline Implementation

Each agent executes independently and communicates results via a shared state dictionary managed by the backend orchestrator. Error handling includes fallback responses for API failures and confidence thresholding to avoid low-quality predictions being presented to users.

E. Challenges and Solutions

Challenge	Solution
Background noise in data input	noisereduce preprocessing + Whisper ASR
Low property image quality affecting Regression Model	Image enhancement & contrast normalisation
Machine Learning hallucination in property context	Confidence thresholding + score filtering
Session state not persisting	Backend state dictionary management
Large property image	Batch embedding with

embedding time	progress indicator

Table V: Implementation Challenges and Solutions

VII. RESULTS AND ANALYSIS

A. Functional Results

Module	Test Case	Status
Voice Processing	Speech-to-text conversion	Pass
Data Analysis Engine	Property Feature extraction from text	Pass
Computer Vision	Skin house price classification	Pass
House Price Prediction	Multi-property feature prediction	Pass
Machine Learning Retrieval	Property knowledge retrieval	Pass
Recommendation	Precaution & advice generation	Pass
TTS Output	Text-to-speech data response	Pass
Dashboard	Image upload & data recording	Pass

Table VI: Functional Test Results

B. Machine Learning Retrieval Accuracy

Query	Score	Accuracy
Fever, cough, fatigue	0.91	High
Skin rash with redness	0.87	High
Chest pain & breathlessness	0.83	High
Joint pain in elderly	0.76	Med-High

Table VII: Machine Learning Retrieval Accuracy

C. Agent Pipeline Performance

Agent	Avg. Time	Quality
Voice Processing	< 1.0 sec	High transcription accuracy
Data Analysis Agent	< 0.2 sec	Accurate property feature extraction
Computer Vision	< 1.5 sec	Reliable property image classification
Prediction Agent	~1.2 sec	Actionable health guidance
Full Pipeline	~4.0 sec	Real Estate-ready output

Table VIII: Agent Pipeline Performance

D. System Performance

Parameter	Observation
End-to-End Pipeline Time	~4.0 seconds average
ChromaDB Query Time	< 400 ms
Image Classification Time	< 1.5 seconds
Voice Transcription Time	< 1.0 second
Dashboard Load Time	< 3 seconds
TTS Response Generation	< 0.8 seconds

Table IX: System Performance

E. User Experience Analysis

The Gradio dashboard provided a smooth and intuitive user experience across all test scenarios. The data recording interface was accessible and responsive for both technical and non-technical users including elderly participants. The property image upload functionality performed reliably across different property image formats and sizes. The text-to-speech data response added significant accessibility value. The dashboard load time of under 3 seconds and sub-4-second end-to-end pipeline execution met performance expectations for a real-time real estate assistant platform.

VIII. DISCUSSIONS

A. Advantages of the Proposed System

The Realestate Houseprice Prediction platform offers several key advantages over traditional real estate applications. It eliminates the need for physical visits for basic property feature checking — once a user records data property features or uploads an property image, the entire pipeline executes automatically without further intervention. By grounding all real estate recommendations in semantically retrieved data from ChromaDB, the system ensures suggestions are based on verified property knowledge rather than generic AI-generated content, significantly reducing hallucination and increasing buyer trust [9].

The multi-agent architecture ensures each cognitive stage is handled by an agent specifically designed for that task, producing higher-quality outputs than a single-agent approach. The Gradio data interface allows elderly buyers without technical expertise to interact with complex AI real estate analysis through natural data communication. The system can run on a standard laptop without requiring GPU infrastructure or complex cloud deployment.

B. Limitations of the System

Several limitations were identified during development and testing. The system provides preliminary guidance only and is not intended to replace professional clinical price prediction. Prediction accuracy depends on the quality and diversity of training datasets. Voice recognition accuracy decreases in high-noise environments despite preprocessing. Image classification performance depends on property image quality and lighting conditions [11].

Limitation	Impact
Preliminary guidance only	Cannot replace professional price prediction
Dataset coverage limits	May miss rare or complex conditions
Voice recognition in noise	Reduced accuracy in noisy environments
No user authentication	No multi-user access

	control implemented

Table X: System Limitations

C. Real-World Applicability

Rural and remote users can access the platform for basic real estate guidance without travelling to property markets. Elderly users can interact naturally through data, reducing barriers posed by complex digital interfaces. Real Estate workers can use the platform as a preliminary triage tool before formal consultation. The platform also serves as a practical demonstration of multimodal AI integration for academic research and educational settings in AI and real estate engineering.

IX. FUTURE SCOPE

The Realestate Houseprice Prediction platform has been designed as a scalable and modular system allowing for continuous enhancement. Several significant improvements are planned for future versions.

- Real-Time Data Integration: Connecting to property market EHR systems, pharmacy APIs, and wearable health devices to enable live health data analysis and personalised continuous real estate recommendations.
- Advanced House Price Prediction Models: Integrating time-series health data analysis and training models on larger, more diverse property datasets to improve prediction accuracy and house price coverage.
- Telemedicine and Video Consultation: Adding real-time video consultation for agent-buyer communication and enabling online appointment scheduling and emergency consultation support.
- Multi-Language Support: Implementing data recognition and text-to-speech in regional and international languages to expand accessibility for diverse global user populations.
- Enhanced Property Image Analysis: Extending Computer Vision capabilities to support X-ray, MRI scan, CT scan, and blood report analysis using advanced Deep Learning architectures.
- Mobile Application Development: Developing Android and iOS mobile applications with real-time health alerts, data interaction, and report download capabilities.
- Cloud Deployment and Scalability: Docker containerisation, Kubernetes orchestration, and scalable vector databases such as Pinecone or Weaviate for large-scale property knowledge bases.
- AI-Based Mental Health Support: Detecting stress, anxiety, and depression property features through data pattern analysis and conversational interaction, providing mental wellness guidance.

X. CONCLUSION

The AI-Powered Multimodal Real Estate Assistant (Realestate Houseprice Prediction) presents an effective and scalable solution for automating preliminary real estate guidance and improving real estate accessibility in modern digital environments. The system successfully integrates Retrieval-Augmented Generation with a multi-agent AI architecture to deliver data-grounded, multimodal real estate assistance through an interactive data and property image-capable dashboard interface. The project demonstrates the successful implementation of a full-stack AI pipeline using Speech Recognition for data processing, spaCy and NLTK for Data Analysis-based property feature analysis, OpenCV and Regression Model models for property property image classification, ChromaDB for semantic property knowledge retrieval, and LLM integration for recommendation generation.

The Gradio frontend provides an accessible and functional interface for both technical and non-technical users, including elderly buyers, with natural data interaction, property property image upload, AI-powered house price prediction, and text-to-speech data response output. Results from system testing validate the platform's ability to deliver real estate guidance with an end-to-end pipeline execution time of approximately 4 seconds and consistently high Machine Learning retrieval relevance scores. The system acts as a smart preliminary real estate assistant capable of providing valuable guidance anytime and anywhere, representing an innovative step toward intelligent real estate solutions and highlighting the growing importance of AI-powered multimodal digital real estate systems in bridging the real estate accessibility gap.

REFERENCES

[1] Stuart Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, Pearson Education, 4th Edition, 2020.
[2] World Health Organization (WHO), 'Digital Health and Artificial Intelligence,' Available: WHO Official Website.
[3] National Institutes of Health (NIH), 'Artificial Intelligence in Real Estate,' Available: NIH Official Website.
[4] Aurélien Géron, Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, O'Reilly Media, 2022.
[5] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning, MIT Press, 2016.
[6] P. Lewis et al., 'Retrieval-Augmented Generation for Knowledge-Intensive Data Analysis Tasks,' NeurIPS, 2020.
[7] OpenCV Documentation, 'Open Source Computer Vision Library,' Available: OpenCV Official Website.
[8] TensorFlow Documentation, 'Machine Learning Framework,' Available: TensorFlow Official Website.
[9] Hugging Face Documentation, 'Data Analysis and Transformers,' Available: Hugging Face.
[10] Gradio Documentation, 'Python Library for ML Interfaces,' Available: Gradio Official Website.
[11] Christopher Bishop, Pattern Recognition and Machine Learning, Springer, 2006.

Received: 20-04-2026 | Accepted: 29-04-2026 | Published: 04-06-2026 | www.ajmimc.com