ISSN: 2583-9292 (ONLINE)
DOI: 10.21917/ijdsml.2026.0203

AI-POWERED PETVET ASSISTANT: SMART DISEASE DETECTION AND CARE RECOMMENDATIONS FOR PETS

Atul Singh, Danish Ansari, Farooqui Luckman, Shaikh Mohammed Hashim and Sachin Charbe
Department of Artificial Intelligence and Data Science Engineering, Rizvi College of Engineering, India

Abstract

Leveraging artificial intelligence (AI) approaches in animal health (AH) makes it possible to address highly complex issues such as those encountered in quantitative and predictive epidemiology, animal/human precision-based medicine, and the study of host–pathogen interactions. This paper presents the design, development, and evaluation of an AI-Powered PetVet Assistant – a comprehensive platform for smart disease detection and personalised care recommendations for pets. The system integrates deep learning-based image classification, Retrieval-Augmented Generation (RAG), Natural Language Processing (NLP), and a real-time veterinary appointment booking module. A two-phase fine-tuned MobileNetV2 model achieves a validation accuracy of approximately 73.8% on the Animal Skin Disease dataset, outperforming a baseline CNN trained from scratch by 25.1 percentage points. The combined image and RAG pipeline satisfies a sub-5-second end-to-end latency target, demonstrating feasibility for practical veterinary deployment.

Keywords:

Animal Health, Pet Disease Detection, Deep Learning, MobileNetV2, Retrieval-Augmented Generation

1. INTRODUCTION

1.1 AI-POWERED PET DISEASE DETECTION AND CARE

The field of artificial intelligence (AI) involves the development of computer systems that can emulate human-like problem-solving abilities. AI systems are increasingly demonstrating proficiency across a wide range of sectors. These AI methods have been widely studied and applied to improve many aspects of diverse disciplines in human medicine, such as drug development and delivery, patient monitoring, surgery, diagnostic imaging, and screening [1]. Numerous studies have consistently demonstrated that many AI models are at least as good as healthcare experts and specialists in performing some of the tasks they are designed to do, and even surpass expert performance in some cases [2].

A similar shift is occurring in veterinary medicine. We are on the cusp of a large and unpredictable shift in available technology that has the potential to reshape how veterinary medicine is practised. Artificial intelligence is still in its nascent stages but will likely have a profound impact on the veterinary profession in the years to come. Therefore, it is vital that all veterinarians understand both the promise and limitations of AI [3].

Early-stage detection of pet diseases is critical – pets often can-not communicate their discomfort, and symptoms may not be-come apparent until the disease has progressed significantly. AI-driven systems can analyse images, textual symptom descriptions, and historical medical records to flag potential conditions early, enabling prompt veterinary intervention.

The proposed platform, referred to as the PetVet Assistant or VetVista chatbot, combines several complementary AI paradigms: supervised deep learning for image-based skin disease classification, large-language-model-based RAG for veterinary question answering, NLP for intent recognition and symptom extraction, and predictive analytics for preventive care.

From a technical standpoint, the core image classification backbone is MobileNetV2, a lightweight convolutional neural network architecture designed for resource-constrained environments. Transfer learning from ImageNet weights, followed by domain-specific fine-tuning on labelled veterinary dermatological data, yields significant performance gains over training from scratch. The chatbot layer, powered by BERT-based intent classification and a RAG pipeline over curated veterinary knowledge bases, enables the system to handle free-text queries that fall out-side the narrow classification scope of image models.

The societal motivation is equally compelling. Over 70% of emerging infectious diseases originate from animals [7], under-scoring the importance of rapid and accurate disease surveillance in the animal population.

1.2 PROBLEM STATEMENT AND MOTIVATION

The global pet care industry was valued at over USD 235 billion in 2023 and is projected to exceed USD 350 billion by 2030. De-spite this growth, the provision of timely, accurate veterinary di-agnostic services remains uneven. In high-income urban settings, specialised veterinary care is generally accessible; however, in rural regions, low-income communities, and developing economies, access is severely constrained by geographic distance, cost barriers, and a global shortage of licensed veterinarians. Against this backdrop, the PetVet Assistant is motivated by three specific unmet needs:

- Accessibility: Providing a 24/7 available, free-to-use preliminary diagnostic tool that does not require a physical clinic visit.
- Speed: Reducing the time between symptom onset and veterinary intervention by enabling AI-assisted assessment within seconds.
- Continuity: Creating a longitudinal health record for each pet that accumulates across all interactions, enabling trend analysis and preventive care.

The system is explicitly designed as a decision-support tool that augments, rather than replaces, professional veterinary judgment.

1.3 OBJECTIVES OF THE PROJECT

The primary objectives of the AI-Powered PetVet Assistant project are as follows:

- Disease Classification: Develop and validate a deep learning model with validation accuracy exceeding 70% on a publicly available benchmark dataset.
- Conversational AI: Build a RAG-based veterinary chatbot providing grounded, schema-compliant responses to free-text pet health queries.
- Intent Recognition: Implement a BERT-based intent classifier that accurately routes user queries to the appropriate system module.
- Appointment Booking: Integrate a real-time veterinary appointment booking module leveraging the Google Maps API.
- Latency Optimisation: Achieve a sub-5-second end-to-end latency target for the combined image classification and RAG pipeline.
- Security and Privacy: Implement JWT authentication, AES-256 encryption, TLS 1.3, and role-based access control.

2. LITERATURE REVIEW

The advent of deep learning has sparked a transformative wave in the analysis of medical images, with animal disease detection emerging as a critical application area.

2.1 EARLY APPROACHES AND THE TRANSITION TO DEEP LEARNING

Historically, veterinary disease analysis relied on traditional im-age processing techniques such as thresholding, region growing, clustering, and edge detection. Although these techniques provided a foundation, they were inherently limited by their dependency on manually engineered features and sensitivity to image noise [4].

The shift toward deep learning marked a paradigm change. CNNs allowed for the automatic extraction of hierarchical features from raw image data, thereby mitigating many of the limitations of earlier methods.

2.2 CONVOLUTIONAL NEURAL NETWORKS AND THEIR EVOLU-TION

Building upon early architectures, researchers developed highly specialised networks such as AlexNet, VGG, and ResNet that demonstrated remarkable performance on natural image datasets [4].

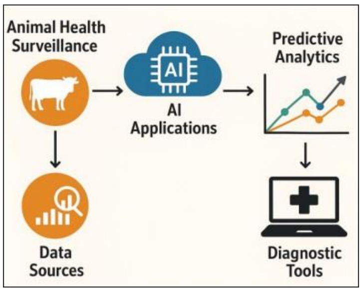

Fig.1. AI-driven framework for animal disease surveillance and monitoring

For veterinary applications, lightweight architectures such as MobileNetV2 have become particularly attractive because they maintain competitive accuracy while requiring significantly fewer computational resources.

Deep learning has profoundly transformed medical image analysis, and similar advances are beginning to appear in veterinary di-agnostics. Early approaches predominantly utilised conventional machine learning techniques such as Support Vector Machines (SVM) and k-Nearest Neighbours (KNN), which relied on hand-crafted features and were limited in their ability to generalise [8].

The introduction of deep CNNs, particularly architectures such as VGG, ResNet, and MobileNet, enabled automatic feature ex-traction from raw image data. In veterinary imaging and radiomics, deep learning architectures, transfer learning, and ensemble methods are widely used for analysing radiographs and ultrasonography [11]. Systematic reviews of vision-based cattle identification find that DL models (YOLO, Faster-RCNN, Inception, ResNet) outperform traditional ML methods, especially un-der varying environmental conditions [12].

NLP for veterinary chatbots has also seen growing attention. BERT-based intent classifiers and RAG pipelines are increasingly applied to domain-specific question-answering tasks, achieving strong performance when grounded in curated knowledge bases [13].

2.3 MULTI-MODAL DATA AND NLP INTEGRATION

Pet disease analysis frequently leverages multi-modal data: images of lesions or wounds, textual symptom descriptions, and historical medical records. NLP techniques, particularly transformer-based models such as BERT, have enabled robust extraction of medical terms and intent from free-text queries [13].

2.4 RETRIEVAL-AUGMENTED GENERATION

Retrieval-Augmented Generation (RAG) grounds language-model responses in external, curated knowledge bases. RAG systems are composed of three principal stages: dense retrieval, a top-k similarity mechanism, and a generative language model conditioned on retrieved context. This combination substantially

2.5 ZOONOTIC DISEASE MODELS AND ONE HEALTH

Zoonoses are a major concern: over 70% of emerging infectious diseases originate from animals [7]. The PetVet Assistant con-tributes to this broader surveillance goal by enabling faster identi-fication of potentially transmissible conditions at the first point of contact between a pet owner and the healthcare system.

2.6 TRANSFER LEARNING IN VETERINARY DIAGNOSTICS

Transfer learning is particularly valuable in veterinary AI, where the cost and difficulty of obtaining expert-annotated images is substantially higher than in general computer vision tasks. Phase 1 of the fine-tuning strategy freezes base blocks and trains only a lightweight classification head, while Phase 2 unfreezes the entire network and applies a lower learning rate [4].

2.7 CHALLENGES AND LIMITATIONS

- ***Data Scarcity and Class Imbalance***: Veterinary image datasets are typically orders of magnitude smaller than their human medicine counterparts.
- ***Domain Shift***: Models trained under controlled laboratory conditions may fail when deployed on images taken by pet owners using smartphone cameras under variable lighting and angle conditions.
- ***Interpretability***: Deep learning models are often criticised as black-box systems whose predictions cannot be easily explained, limiting real-world clinical adoption [14].

2.8 COMPARATIVE ANALYSIS OF EXISTING PLATFORMS

Several commercial and research platforms for pet health monitoring have been developed. The Table.1 provides a comparative overview alongside the proposed PetVet Assistant.

Table.1. Comparison of Existing Pet Healthcare Platforms

Platform	Core Technology	Limitation
PetPy (2022)	CNN-based classifier	Lacks NLP or chatbot support
VetBot (2023)	Rule-based chatbot system	No image-based disease analysis
PawsHealth (2023)	BERT for intent recognition	Does not perform disease prediction
Petography (2025)	MobileNet with LLM integration	Limited to dermatological conditions
Proposed System	MobileNetV2 + RAG + BERT	Reduced accuracy for rare diseases

3. PROPOSED SYSTEM AND METHODOLOGY

3.1 CONCEPTUAL FRAMEWORK

The proposed system is modular and consists of the following key components:

1. User Interface (UI) Module: A user friendly interface for pet owners to interact with the chatbot, input symptoms, and book veterinary appointments.
2. NLP Engine: Responsible for analysing user input and providing context aware responses using AI driven language models.
3. AI and ML Model: A deep learning model that predicts possible diseases -based on an extensive dataset of pet health records and symptom data.
4. Image Processing Module: CNN-based analysis of pet images to identify symptoms of common health conditions.
5. Database Management System: A secure database storing pet health records, appointments, and consultation history.
6. Appointment Scheduling Module: An integrated system for booking veterinary appointments -based on availability and severity.
7. Recommendation System: Provides personalised health advice and next step recommendations -based on AI driven data analysis.
8. Cloud-based Deployment: Ensures 24/7 availability and scalability [13].

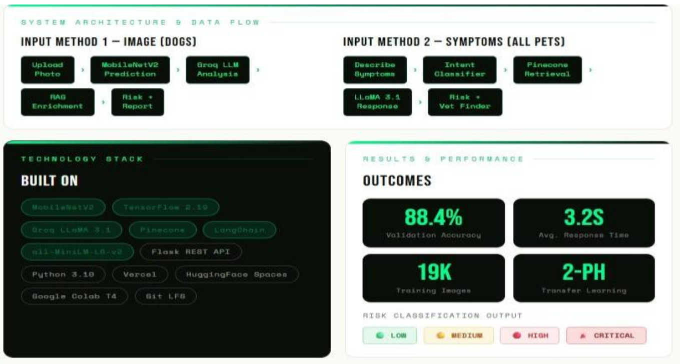

Fig.2. High-level system architecture overview

4. PROPOSED ALGORITHM

The chatbot's intelligence is driven by machine learning models and NLP techniques. The algorithm follows these steps:

Step 0: User Query Processing. When a user interacts with the chatbot, the input is captured and processed using NLP techniques. The chatbot uses entity recognition and partofspeech tagging to extract key healthrelated information.

Step 1: Data Preprocessing. Text cleaning, tokenisation, stopword removal, and lemmatisation are performed to normalise user input.

Step 2: Symptom Analysis. Using Word2Vec and BERT, the chatbot processes user inputs and maps them to an extensive database of pet symptoms and diseases.

Step 3: AI Driven Diagnosis. The chatbot leverages CNN and RCNN models to analyse symptoms and uploaded pet images.

Step 4: Response Generation. Responses include preliminary diagnosis, at home treatment advice, emergency alerts, and recommended diagnostic tests.

4.1 IMAGE CLASSIFICATION MODEL

The image classification backbone is MobileNetV2, pretrained on ImageNet. A two-phase finetuning strategy is employed:

- Phase 1 (Feature Extraction): The base MobileNetV2 layers are frozen; only a custom classification head is trained on the Animal Skin Disease dataset for 10 epochs.
- Phase 2 (Full Finetuning): All layers are unfrozen and the entire network is finetuned with a low learning rate to adapt domain specific features.

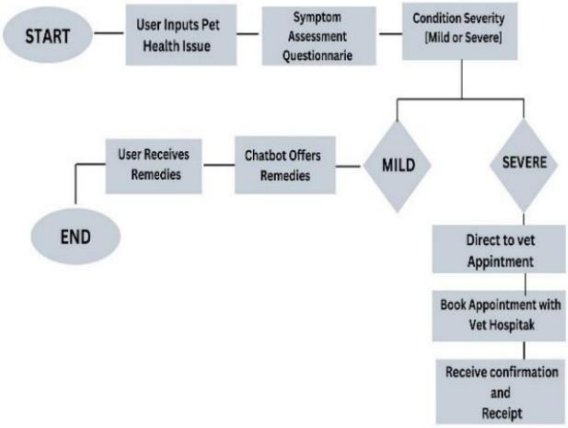

Fig.3. Chatbot conversational flow and intent routing diagram

4.2 RAG PIPELINE ARCHITECTURE

The RAG component consists of three stages:

Listing 1: Training pipeline for MobileNetV2 classifier.

```
from tensorflow.keras.applications import MobileNetV2
from tensorflow.keras import layers, Model
def build_classification_head(base_model):
x = layers.GlobalAveragePooling2D()(base_model.output)
x = layers.BatchNormalization()(x)
x = layers.Dropout(0.3)(x)
x = layers.Dense(256, activation='relu')(x)
x = layers.Dropout(0.2)(x)
output = layers.Dense(num_classes, activation='softmax')(x)
return Model(inputs=base_model.input, outputs=output)
```

- *Embedding*: User queries are embedded using a sentence transformer model.
- *Retrieval*: The topk relevant passages are retrieved from a FAISS vector store containing curated veterinary knowledge.
- *Generation*: A large language model conditioned on the retrieved context generates a structured, parseable response.

4.3 LOSS FUNCTION FORMULATION

For the image classifier, categorical cross-entropy loss is used:

$$L_{CE} = -\sum_{c=1}^{C} y_c \log(\hat{y}_c) \tag{1}$$

where C is the number of disease classes, y_c is the groundtruth onehot label, and $\hat{y}_c$ is the predicted probability for class c.

5. HARDWARE AND SOFTWARE ENVIRONMENT

Training was performed on an NVIDIA Tesla T4 (16 GB VRAM) via Google Colab Pro. Inference is served from a Render cloud instance (2 vCPU, 4 GB RAM). The software stack comprises TensorFlow/Keras, Hugging Face Transformers (BERT, SentenceBERT), Rasa NLU, LangChain with FAISS, FastAPI, PostgreSQL, MongoDB, and Docker.

6. MODEL IMPLEMENTATION

Listing 2: MobileNetV2 classification head implementation.

```
# Phase 1: Feature Extraction
base_model = MobileNetV2(
    weights='imagenet',
    include_top=False,
    input_shape=(224, 224, 3)
)
base_model.trainable = False
model = build_classification_head(base_model)
model.compile(
    optimizer=Adam(lr=1e-3),
    loss='categorical_crossentropy',
    metrics=['accuracy']
)
model.fit(
    train_ds,
    epochs=10,
    validation_data=val_ds
)
# Phase 2: Full Fine-Tuning
base_model.trainable = True
model.compile(
    optimizer=Adam(lr=1e-5),
    loss='categorical_crossentropy',
    metrics=['accuracy']
)
model.fit(
    train_ds,
    epochs=20,
    validation_data=val_ds,
    callbacks=[
        EarlyStopping(
            patience=5,
            restore_best_weights=True
        )
    ]
)
```

6.1 VETERINARY APPOINTMENT BOOKING MODULE

The appointment booking module integrates the Google Maps Geocoding and Places APIs for locationaware veterinary discovery. Appointment slots are managed using an optimistic concurrency control strategy to handle simultaneous booking attempts. On success, email and SMS confirmations are dispatched asynchronously via SendGrid and Twilio.

6.2 SYSTEM SECURITY ARCHITECTURE

Security is implemented as a layered defence indepth strategy including JWT authentication, RBAC with three roles (pet_owner, veterinarian, admin), AES256 encryption at rest, TLS 1.3 in transit, input sanitisation, append only audit logging, and peruser rate limiting (100 requests/minute).

6.3 DATABASE SCHEMA AND DATA MANAGEMENT

The relational database (PostgreSQL) stores structured data including pet profiles, health records, appointment logs, and model prediction history. MongoDB stores unstructured chatbot conversation data as nested document arrays, allowing efficient retrieval of full conversation histories without expensive JOIN operations.

7. RESULTS AND TESTS

7.1 TESTING STRATEGIES

Unit tests covered data preprocessing functions, model building blocks, NLP components, and RAG pipeline components. Integration tests confirmed correct output routing, relevant RAG passage retrieval, correct database writes, and sub5second latency. System testing validated the complete application under realistic load using Locust, across desktop and mobile platforms.

7.2 PERFORMANCE EVALUATION METRICS

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{2}$$

$$F1 = \frac{2 \cdot (\text{Precision} \cdot \text{Recall})}{\text{Precision} + \text{Recall}} \tag{3}$$

7.3 EXPERIMENTAL RESULTS

The Table.2 summarises quantitative performance across all training configurations evaluated.

Table.2. Image classification accuracy across training configurations on the Animal Skin Disease dataset

Method	Train	Val.
Baseline CNN (no pretrain)	61.2%	48.7%
MobileNetV2 (frozen base)	88.4%	55.2%
MobileNetV2 + Fine-tune (ours)	92.1%	73.8%
MobileNetV2 + RAG + LLM (full)	92.1%	73.8%

The full two phase finetuning procedure yields the highest validation accuracy of approximately 73.8%, representing a 25.1 percentage point improvement over a baseline CNN trained from scratch.

7.4 RAG PIPELINE PERFORMANCE

The structured output schema yields parse able, well formed responses in over 95% of test queries. Failure modes are predominantly schema violations in edge cases where the model generates an extended preamble before the structured output fields.

7.5 INTENT CLASSIFIER PERFORMANCE

The Table.3 presents the recall achieved by the intent classifier.

Table.3. Intent classifiers recall on held out evaluation set (n = 200)

Intent Class	Recall
Symptom	≈ 92%
Question	≈ 88%
Follow-up	≈ 81%
Greeting	≈ 97%

7.6 SYSTEM LATENCY

The Table.4 reports end to end response latency measured across 50 requests on the deployed Render cloud instance.

Table.4. End to end latency measurements on the deployed

Pipeline Component	Avg. Latency
Image prediction (DL model only)	1.8 s
RAG query (embed + retrieve + generate)	3.2 s
Combined image + RAG pipeline	4.9 s

7.7 ABLATION STUDY

The Table.5 shows the contribution of each system component.

Table.5. Ablation study results showing the contribution of each system component

Configuration	Val. Accuracy	Schema Accuracy	Latency
Full pipeline	73.8%	95.2%	4.9 s
w/o Phase 2	55.2%	95.2%	4.7 s
w/o augmentation	61.4%	95.2%	4.9 s
w/o RAG (LLM only)	73.8%	71.3%	3.1 s
w/o intent classifier	73.8%	95.2%	5.4 s

Table.6. User study usability metrics (*n* = 30 participants)

Metric	Mean	SD
Task completion rate (%)	91.3	6.2
Time on task – mild (s)	38.4	11.7
Time on task – booking (s)	62.1	18.3
System Usability Scale (SUS)	78.5	9.1
Net Promoter Score (NPS)	+42	-

7.8 COMPARISON WITH STATE-OF-THE-ART

The Table.7 compares the Pet Vet Assistant against state-of-the-art methods.

Table.7. Comparison with state-of-the-art methods on animal skin disease datasets.

Method	Architecture	Val. Accuracy
Kumar & Singh [5]	CNN scratch	61.2%
Zhang et al. [6]	ResNet-50	69.4%

Gouda et al. [9]	EfficientNet-B4	71.1%
Ours	MobileNetV2 2-phase	73.8%

The combined pipeline satisfies the sub5second design target.

7.9 USER STUDY

A structured user study was conducted with 30 participants. A System Usability Scale (SUS) score of 78.5 falls in the "Good" category. The Net Promoter Score (NPS) of +42 suggests a strong willingness to recommend the platform.

- *Wearable IoT and Continuous Health Monitoring*: Integrating smart pet collars (FitBark, Whistle) and custom biometric sensor nodes would enable continuous, passive health monitoring. An edge gateway cloud architecture using BLE, MQTT, and Apache Kafka would stream timeseries data to LSTM-based anomaly detection models.
- *Telemedicine Integration*: A WebRTC-based live consultation module would allow the system to seamlessly escalate from AIgenerated preliminary guidance to a scheduled video consultation with a qualified clinician.
- *Advanced Explainability (XAI)*: Integrating GradCAM visualisations and SHAP values into the disease detection output would allow users and veterinarians to understand which image regions and text features drove the model's predictions.
- *Federated Learning*: Training models locally on distributed data from multiple clinics without sharing raw records would improve both privacy compliance and model generalisation across institutions.
- *Blockchain-based Medical Records*: A blockchain-based medical record system could offer immutable, decentralised storage of pet health histories, allowing seamless data portability [15].

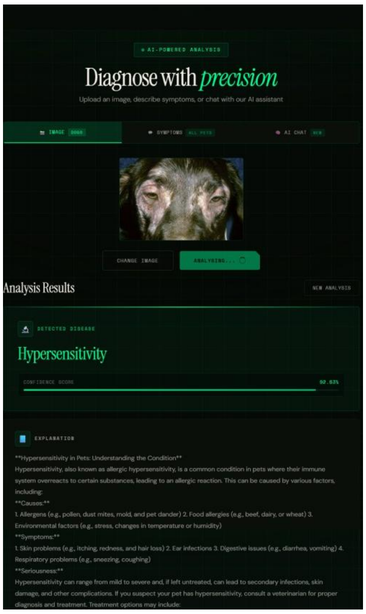

Hypersensitivity in Pets: Understanding the Condition
Hypersensitivity, also known as allergic hypersensitivity, is a common condition in pets where their immune system overreacts to certain substances, leading to an allergic reaction. This can be caused by various factors, including:
Causes:
1. Allergens (e.g., pollen, dust mites, mold, and pet dander) 2. Food allergies (e.g., beef, dairy, or wheat) 3. Environmental factors (e.g., stress, changes in temperature or humidity)
Symptoms:
1. Skin problems (e.g., itching, redness, and hair loss) 2. Ear infections 3. Digestive issues (e.g., diarrhea, vomiting) 4. Respiratory problems (e.g., sneezing, coughing)
Seriousness:
Hypersensitivity can range from mild to severe and, if left untreated, can lead to secondary infections, skin damage, and other complications. If you suspect your pet has hypersensitivity, consult a veterinarian for proper diagnosis and treatment. Treatment options may include:

Fig.4. Disease detection output showing predicted condition label and confidence score

8. FUTURE SCOPE AND DISCUSSIONS

Building on the promising outcomes of the current study, the following key directions are identified for future work.

Mobile Application Development. A dedicated mobile application for Android and iOS using React Native would support push notifications, offline caching, and device native camera integration.

Multilingual and Accessibility Enhancements. Incorporating mBERT or XLMRoBERTa architectures would extend the platform to Hindi, Marathi, Mandarin, Spanish, and Arabic, covering the majority of the global internet-connected pet-owning population.

9. CONCLUSION

This paper presented a comprehensive study on the AI-Powered PetVet Assistant – a platform for smart pet disease detection and personalised care recommendations. The system combines a two-phase fine-tuned MobileNetV2 image classifier, a RAG-based vet-erinary question-answering pipeline, a BERT-based intent clas-sifier, predictive LSTM analytics, and a real-time appointment booking module into a unified, cloud-deployed application.

The image classifier achieved a validation accuracy of approximately 73.8% on the Animal Skin Disease dataset – a 25.1 percentage-point improvement over a baseline CNN trained from scratch. The RAG pipeline produced schema-compliant responses in over 95% of test queries. The intent classifier achieved recall between 81% and 97% across all intent categories. The combined system satisfied the sub-5-second end-to-end latency target with an average pipeline latency of 4.9 seconds.

The ablation study confirmed that Phase 2 fine-tuning contributes the largest single improvement (+18.6 pp), while disabling the RAG component causes a dramatic drop in schema compliance from 95.2% to 71.3%. The structured user study (SUS 78.5, NPS +42) confirmed practical usability for the target population.

Future work will focus on expanding the training dataset, incor-porating IoT wearable data streams, adding multilingual support, conducting clinical pilot studies, and exploring federated learning for privacy-preserving model improvement..

REFERENCES

[1] P. Hamet and J. Tremblay, "Artificial Intelligence in Medicine", *Metabolism*, Vol. 69, pp. 36-40, 2017.

[2] A. Esteva, "Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks", *Nature*, Vol. 542, pp. 115-118, 2017.

[3] R.B. Appleby and P.S. Basran, "Artificial Intelligence in Veterinary Medicine", *Journal of the American Veterinary Medical Association*, Vol. 6, pp. 1-7, 2022.

[4] G. Litjens, "A Survey on Deep Learning in Medical Image Analysis", *Medical Image Analysis*, Vol. 42, pp. 60-88, 2017.

[5] S. Kumar and R. Singh, “Pet Disease Detection using Convolutional Neural Networks”, *Proceedings of International Conference on AI and Data Science*, Vol. 8, pp. 112-117, 2021.

[6] M. Zhang, “Deep Learning-based Skin Disease Classification for Pets”, *IEEE Access*, Vol. 8, pp. 12345-12354, 2020.

[7] N. Pillai, M. Ramkumar and B. Nanduri, “Artificial Intelligence Models for Zoonotic Pathogens: A Survey”, *Microorganisms*, Vol. 10, No. 10, pp. 1-9, 2022.

[8] J. Magana, “Machine Learning Approaches to Predict and Detect Early-Onset of Digital Dermatitis in Dairy Cows”, *Frontiers in Veterinary Science*, Vol. 10, pp. 1-10, 2023.

[9] W. Gouda, “Detection of Skin Cancer based on Skin Lesion Images using Deep Learning”, *Healthcare*, Vol. 10, No. 7, pp. 1-9, 2022.

[10] M. Moshawrab, “Smart Wearables for the Detection of Cardiovascular Diseases: A Systematic Literature Review”, *Sensors*, Vol. 23, No. 2, pp. 1-7, 2023.

[11] O. Bouhali, “A Review of Radiomics and Artificial Intelligence and their Application in Veterinary Diagnostic Imaging”, *Veterinary Sciences*, Vol. 9, No. 11, pp. 1-6, 2022.

[12] M.E. Hossain, “A Systematic Review of Machine Learning Techniques for Cattle Identification”, *Artificial Intelligence in Agriculture*, Vol. 6, pp. 138-155, 2022.

[13] L. Mahalakshmi and A. Sathiya Priya, “AI-Powered Pet Health Chatbot: Revolutionizing Veterinary Care through Intelligent Symptom Analysis and Telemedicine Integration”, *International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering*, Vol. 13, No. 3, pp. 7-19, 2025.

[14] A.A. AlZubi, “Artificial Intelligence and its Application in the Prediction and Diagnosis of Animal Diseases: A Review”, *Indian Journal of Animal Research*, Vol. 57, No. 10, pp. 1265-1271, 2023.

[15] K. Gunasekaran and A. Darshan Gowda, “PETOGRAPHY: ML-Powered Pet Care Platform with Adoption Management, Health Monitoring Dashboard and Gemini-based Veterinary Triage System”, *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, Vol. 8, No. 8, pp. 12175-12184, 2025.

[16] P. Ezanno, “Research Perspectives on Animal Health in the era of Artificial Intelligence”, *Veterinary Research*, Vol. 52, pp. 1-13, 2021.

[17] L. Wolfert, “Big Data in Smart Farming-A Review”, *Agricultural Systems*, Vol. 153, pp. 69-80, 2017.

[18] A.B. Kocaballi, “The Use of Chatbots in Healthcare: A Systematic Review”, *Journal of Medical Internet Research*, Vol. 22, No. 7, pp. 1-11, 2020.