IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

MONKEY POX DISEASE DETECTION USING DEEP LEARNING

R. Madhananthesh , Dr. Kamlesh Tiwari Department of Computer Science and Engineering (DS)

Jain University Bengaluru, India

Abstract

The recent global re-emergence of Monkeypox (Mpox) necessitates rapid, non- invasive diagnostic tools to differentiate its dermatological manifestations from visually similar diseases such as Chickenpox and Measles. While deep learning models possess high theoretical accuracy for skin lesion classification, standard architectures frequently fail in clinical deployment due to "black-box" opacity and their structural vulnerability to out-of-distribution (OOD) inputs—often generating high-confidence false-positive diagnoses when presented with non-medical imagery. In this paper, we propose a robust, multi-stage Explainable Artificial Intelligence (XAI) framework to resolve these limitations. The core mathematical feature extractor utilizes an EfficientNetB0 architecture trained across five distinct classes, including a dynamic "Others" category to compartmentalize generic skin anomalies. To enforce diagnostic safety, a lightweight MobileNetV2 semantic sentry is integrated to instantaneously reject OOD inputs (such as animals or household objects) prior to inference, alongside an 80% algorithmic confidence threshold. Finally, the framework employs Gradient-weighted Class Activation Mapping (Grad-CAM) to visually project the specific convolutional features influencing the model's prediction, thereby enforcing clinical accountability. By combining high-fidelity transfer learning with explicit OOD defense mechanisms and real-time visual explainability, our proposed system bridges the critical gap between deep learning theoretical accuracy and practical, trustworthy medical deployment.

Keywords:

Monkeypox Detection, Deep Learning, EfficientNetB0, Explainable AI (XAI), Grad- CAM, Out-of-Distribution (OOD) Detection, Computer-Aided Diagnosis (CAD).

1. Introduction

Monkeypox (Mpox), a viral zoonosis caused by the Orthopoxvirus, has historically remained endemic to Central and West Africa. However, recent unprecedented global outbreaks have rapidly elevated the disease to a Public Health Emergency of International Concern. The primary vector for clinical diagnosis is the manifestation of distinct dermatological eruptions, which sequentially progress from macules to fluid-filled vesicles and eventually scabs. A critical bottleneck in outbreak containment is the initial diagnosis; the visual progression of Monkeypox lesions is notoriously difficult to differentiate from other widespread exanthematous diseases such as Chickenpox and Measles, often confounding practitioners without immediate

access to laboratory Polymerase Chain Reaction (PCR) testing.

In recent years, Artificial Intelligence—specifically through Convolutional Neural Networks (CNNs)—has demonstrated exceptional capabilities in the realm of computer-aided diagnosis (CAD) for dermatology. Deep learning topologies possess mathematical proficiencies for extracting and mapping microscopic textural and structural variations in skin lesions that are frequently imperceptible to the human eye. Despite achieving high theoretical accuracy on curated datasets, the widespread clinical adoption of these deep learning models remains heavily bottlenecked by two persistent systemic flaws.

First is the "black-box" paradigm. Traditional CNN architectures yield discrete probabilistic outputs without explaining the internal visual features driving the conclusion. Within the stringent ethical confines of healthcare, diagnostic tools that lack transparent reasoning are untrustworthy and ultimately rejected by clinicians. Second is the profound vulnerability of standard models to Out-of-Distribution (OOD) inputs. AI models bounded tightly within specific medical datasets inherently lack the semantic awareness to identify non-medical noise; consequently, when a user uploads irrelevant imagery (such as a household object or a pet), conventional models are mathematically forced to classify the noise as a medical condition, resulting in catastrophic false-positive clinical assertions.

To address these critical limitations, this paper proposes an end-to-end, highly defensive Explainable AI (XAI) framework optimized for real-world triage environments. The specific contributions of this paper are as follows:

1. Multi-Class Feature Extraction: The implementation of an EfficientNetB0 architecture leveraging a 5-class dataset structure—including a definitive "Others" category to safely sequester anomalous, non-target dermatological conditions.
2. Algorithmic Out-of-Distribution Defense: The integration of a pre-inference semantic sentry module powered by MobileNetV2. This module actively blocks and redirects gross visual anomalies (e.g., animals whose fur texture mimics skin) before medical analysis can be erroneously executed.
3. Clinical Transparency: The active deployment of Gradient-weighted Class Activation Mapping (Grad-CAM) to generate real-time algorithmic heatmaps, explicitly localizing the physical lesion characteristics that precipitated the network's diagnostic decision.
4. Predictive Safety Constraints: The enforcement of mathematical boundaries, explicitly neutralizing low-confidence predictions (<80%) as "Inconclusive" to maintain strict diagnostic integrity.

2. Related Work

The application of deep learning algorithms to dermatological imaging has been extensively explored, yielding significant advancements in automated disease classification. However, the specific diagnosis of Monkeypox via algorithmic triage is a relatively nascent field, and the existing body of literature predominantly suffers from distinct architectural and operational limitations that this paper seeks to resolve.

2.1 Conventional Convolutional Neural Networks in Dermatology

Initial research in automated skin lesion classification heavily utilized foundational convolutional topologies such as VGG16, InceptionV3, and ResNet-50. For instance, early diagnostic frameworks for melanoma and exanthematous conditions demonstrated that transfer learning on ImageNet-initialized weights could yield diagnostic accuracy comparable to human dermatologists. However, these legacy architectures are notoriously over-parameterized. Deploying models like VGG16 or ResNet-50 in real-world, latency-sensitive web environments introduces severe computational overhead. Recent state-of-the-art literature has initiated a shift toward optimized networks. The introduction of EfficientNet paradigms—which balance network depth, width, and resolution scaling—has proven capable of extracting high-fidelity pathological features with a fraction of the computational burden. Consequently, this study

adopts EfficientNetB0 as the core mathematical extractor to optimize inference efficiency without sacrificing diagnostic integrity.

2.2 The Oversight of Out-of-Distribution (OOD) Fallibility

A catastrophic flaw pervading the majority of current diagnostic CAD literature is the inherent assumption of strictly homogeneous input data. Diagnostic models are predominantly trained, validated, and tested within closed ecosystems of pure medical imagery. Consequently, standard literature rarely addresses Out-of- Distribution (OOD) vulnerability. When standard models are deployed publicly, the ingestion of non-medical noise (e.g., household objects, wildlife, or varied human profiles) forces the algorithm to mathematically map the anomaly to a disease class, generating dangerous false positives. A handful of contemporary studies have attempted algorithmic rejection using basic color-chrominance thresholding (such as HSV or YCrCb skin-color detection). However, these heuristics are fundamentally brittle, frequently resulting in false negatives when presented with pale skin, dark skin, or dense hair. This study diverges significantly from existing literature by proposing an active, semantic MobileNetV2 gatekeeper strategy designed to algorithmically filter anomalous objects based on ImageNet topological features, entirely bypassing the vulnerabilities of rigid color thresholding.

2.3 The Necessity of Explainable AI (XAI)

Furthermore, a substantial portion of existing Monkeypox classification studies focus exclusively on optimizing raw probabilistic metrics (e.g., F1-Scores and Top-1 Accuracy) while ignoring the "Black-Box" opacity of the underlying neural networks. In formal clinical settings, opaque algorithmic outputs are frequently deemed legally and ethically problematic. Recent diagnostic frameworks have begun integrating Explainable AI (XAI) to foster clinical trust. By utilizing Gradient- weighted Class Activation Mapping (Grad-CAM), recent methodologies can interpolate exactly which spatial textures influenced the decision layer. Building upon this, our proposed framework tightly couples Grad-CAM to the `top_conv` layer of the EfficientNet base, ensuring that every inference is accompanied by a transparent, visually interpretable heatmap overlay.

3. Proposed System

The proposed framework is an end-to-end diagnostic pipeline designed to ensure robustness, accuracy, and interpretability through a structured multi-stage process.

3.1 Out-of-Distribution (OOD) Validation

To prevent non-medical inputs, a lightweight MobileNetV2 classifier performs initial screening. Inputs identified as irrelevant (e.g., animals or noise) are rejected, ensuring only medically valid images proceed for diagnosis.

3.2 Preprocessing and Feature Extraction

Validated images are resized to 224 × 224 × 3 and normalized. Feature extraction is performed using EfficientNetB0, leveraging transfer learning. Initially, base layers are frozen, followed by gradual fine-tuning to adapt the model for dermatological feature representation.

3.3 Classification Head

Extracted features are processed using Global Average Pooling and dense layers (1024, 512 units) with Dropout (0.5, 0.3) to reduce overfitting. A Softmax layer outputs probabilities across five classes: Monkeypox, Chickenpox, Measles, Normal Skin, and Others.

3.4 Safety Constraint and Explainability

A minimum confidence threshold of 80% is applied; predictions below this are labeled Inconclusive. Model decisions are visualized using Grad-CAM, generating heatmaps to highlight diagnostically relevant regions and enhance interpretability.

3.1. System Architecture

The proposed system is a sequential diagnostic pipeline designed to filter invalid inputs and extract clinically relevant features for accurate skin disease classification. It consists of four main stages:

3.1 Input Acquisition and OOD Validation

The pipeline begins with a user-uploaded skin lesion image. An OOD validation module based on MobileNetV2 screens the input to detect non-medical images (e.g., animals). Invalid inputs are rejected, while valid clinical images proceed for further analysis.

3.2 Preprocessing and Feature Extraction

Validated images are resized to 224 × 224 and normalized. Feature extraction is performed using EfficientNetB0 with transfer learning. Initial layers are frozen to retain general features, while top layers are fine-tuned to capture disease-specific patterns.

3.3 Classification Head

Extracted features are processed via Global Average Pooling and dense layers (with Dropout for regularization). The final layer outputs class probabilities for multiple skin conditions using Softmax.

3.4 Frontend Output and Explainability

Results are displayed through a web interface, providing:

Predicted Class with Confidence Score

Grad-CAM Heatmap, highlighting important regions influencing the model's decision for improved interpretability.

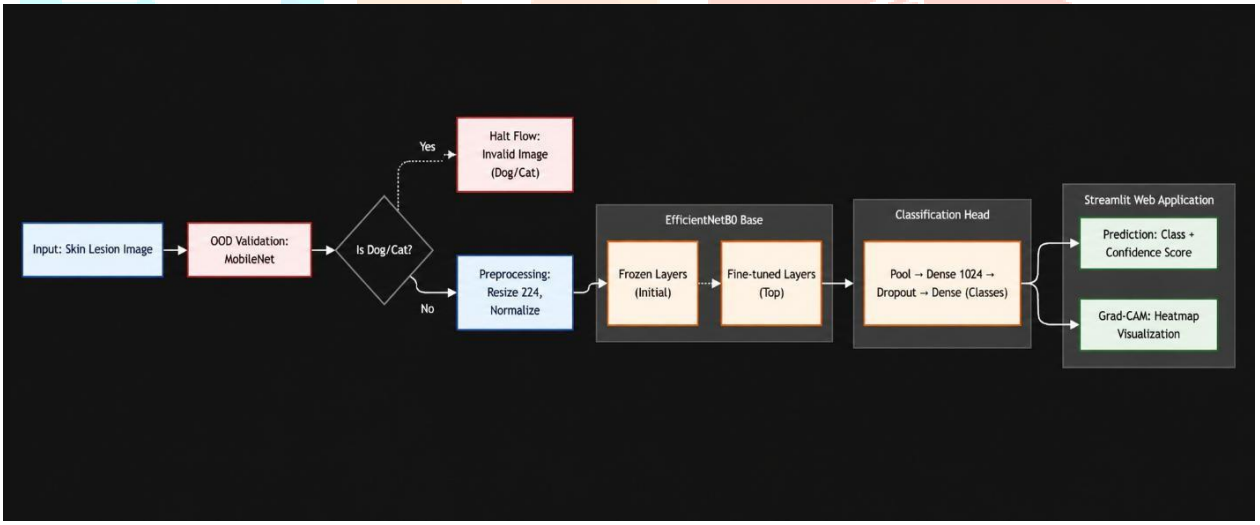

Figure 1: System architecture of AI Wellness Companion

4. Materials and Methods

This section systematically delineates the fundamental datasets, preprocessing techniques, algorithmic strategies, and environmental configurations utilized to construct, train, and mathematically validate the proposed deep learning framework.

4.1 Materials: Dataset Acquisition and Segmentation

The foundational material for this study comprises a highly curated repository of clinical dermatological imagery. Recognizing the severe limitations of binary classification (which forces the misclassification of unmapped diseases), researchers constructed an explicit 5- class diagnostic dataset. The repository is categorized into:

1. Monkeypox (Umbilicated papules, severe vesicles, and scabbing stages).
2. Chickenpox (Superficial, distinct clustered vesicles).

3. Measles (Maculopapular erythematous rashes).
4. Normal Skin (Healthy epidermal control data).
5. Others (A deliberate compilation of miscellaneous dermatological anomalies such as eczema, generic acne, and benign nevi).

The inclusion of the `Others` category is a critical methodological safeguard, training the neural network to actively sequester generalized skin noise rather than forcing a false- positive clinical match. All images were digitally standardized to an absolute spatial resolution of $224 \times 224 \times 3$ to comply with the rigid input tensor constraints of the core classification network.

4.2 Data Augmentation Methodology

To artificially expand the variance of the training dataset and aggressively mitigate the risk of algorithmic memorization (overfitting), a dynamic data augmentation pipeline was deployed using Keras `ImageDataGenerator`.

During the training phase, tensors were subjected to randomized geometric perturbations. Operations included stochastic horizontal flipping, deliberate shear mapping, minor zoom fluctuations, and arbitrary rotational shifts up to bounds of 20 degrees. This methodology mathematically forces the deep learning topology to identify scale-invariant and rotationally-agnostic pathological micro-features rather than relying heavily on the default spatial orientation of a captured photograph.

4.3 Transfer Learning and Deep Features Construction

Training a deep network entirely from scratch on a constrained clinical dataset is computationally inefficient and mathematically unstable. Consequently, this study employed a Transfer Learning methodology layered upon the EfficientNetB0 architecture. The model was strategically initialized utilizing pre-trained weights derived from the vast ImageNet repository.

A calculated Two-Stage Fine-Tuning methodology was implemented to preserve these baseline architectural weights:

- Phase I (Architectural Freezing): The entire EfficientNet convolutional base was algebraically frozen. Only the newly appended premium Classification Head—comprising successive Global Average Pooling, highly regularized Dense networks (1024 and 512 neurons respectively), and aggressive Dropout mechanics (0.5 and 0.3)—was allowed to train. This allowed the chaotic, randomly initialized gradient weights of the new head to naturally stabilize without heavily backpropagating noise into the foundational edge-detection blocks.
- Phase II (Deep Feature Unfreezing): The terminal convolutional blocks of the base architecture were subsequently unfrozen. The network was then re-trained leveraging exceptionally low learning rate parameters (e.g., 1×10^{-5}). This gentle phase permitted the network to delicately re-map its deep abstraction layers to specifically correlate with microscopic, disease-specific textural variances (such as the physiological depth of a Monkeypox lesion).

4.4 Experiential Setup and Optimization Metrics

During network convergence, algorithmic feedback was mathematically measured utilizing a standard Categorical Cross-Entropy loss function. To traverse the loss landscape efficiently, the Adam (Adaptive Moment Estimation) optimization algorithm was strategically implemented over the default Stochastic Gradient Descent. The network's probability distribution constraint was enforced via an ultimate 5-node Softmax topology.

4.5 Environmental and Software Configuration

The overarching computational framework was strictly developed utilizing mature open- source topologies. Python 3 constituted the core mathematical language. Deep learning model fabrication, layer architecture definition, and gradient descents were managed via the TensorFlow 2 / Keras API structural backbone. Real-time geometric image manipulation was supported by OpenCV, and the final algorithmic logic was rendered for public triage access via the Streamlit cloud framework.

5. Implementation Workflow

The physical realization of the proposed diagnostic architecture adhered to a strict, multi- phase developmental and operational pipeline. The workflow sequence ensures that mathematical models trained in isolated ecosystems are robust enough to transition cleanly into real-time clinical triage software. The deployment pipeline was abstracted into five distinct phases:

Phase 1: Data Structuring and Environmental Initialization

The workflow initiated with the construction of the underlying computational environment utilizing Python 3 and the localized installation of specific TensorFlow/Keras and OpenCV libraries. Following environment initialization, the raw multi-source image data was mathematically consolidated into five explicit categorical directories (Monkeypox, Measles, Chickenpox, Normal, and Others). A programmatic data-splitting script geometrically segregated the repository into an 80:20 training-to- validation ratio, ensuring that algorithmic optimization was entirely independent of the foundational testing subsets to brutally evaluate generalization capabilities.

Phase 2: Network Topology Configuration and Parameter Binding

Following data stratification, the central algorithmic extraction core was initialized. The `EfficientNetB0` architectural topology was computationally instantiated utilizing base ImageNet weight configurations. To transition the model from generalized semantic abstraction to dedicated dermatological profiling, a custom analytical head was appended. Specifically, a Global Average Pooling 2D tensor reduction was mapped to dual dense strata (1024 and 512 neurons). Aggressive spatial dropout variables (at constraints of 0.5 and 0.3) were injected between these strata to mathematically penalize gradient over-reliance during backpropagation.

Phase 3: Transfer Optimization and Epoch Iteration

The constructed topology underwent rigorous compilation via the `Adam` optimizer bounded by a Categorical Cross-Entropy objective function. The training iteration loop mathematically adopted a Two-Stage Fine-Tuning paradigm. In the preliminary epochs, the base EfficientNet was fully restricted (frozen) to allow the customized classification stratus to stabilize its random geometric vectors. Upon stabilization, the terminal blocks of the EfficientNet foundation were programmatically unfrozen, and the model resumed epoch iterations constrained by a severely restricted learning rate schedule (1×10^{-5}), ensuring high-fidelity convergence upon microscopic exanthematous variances.

Phase 4: Out-of-Distribution (OOD) Sentry Calibration

To mathematically insulate the newly trained classifier from catastrophic external failure upon deployment, a secondary, lightweight topology—MobileNetV2—was introduced into the operational pipeline. It was programmed as an active semantic sentry. Specifically, custom filtering algorithms were encoded so that if the MobileNet topology yielded high-confidence predictions mapping to internal ImageNet animal classes (indexing keys 151 through 285) or standard objects, an algebraic short-circuit override would immediately trigger, safely terminating the primary medical analysis route.

Phase 5: Explainable Integration and Cloud Application Deployment

The final workflow phase transitioned raw Python logic into the clinical presentation layer utilizing the Streamlit cloud topology. The underlying mathematics were encapsulated within a functional frontend interface. Concurrently, a parallel Grad-CAM extraction function was bound algebraically to the final computational block of the EfficientNet model (`top_conv`). During a live clinical query, if the Softmax logit exceeded the stringent 80% confidence parameter, the probability array triggers the Grad-CAM block. The resulting spatial gradients generate an independent visual heatmap map structure, which is subsequently warped and seamlessly overlaid onto the user's primary source image within the Streamlit User Interface.

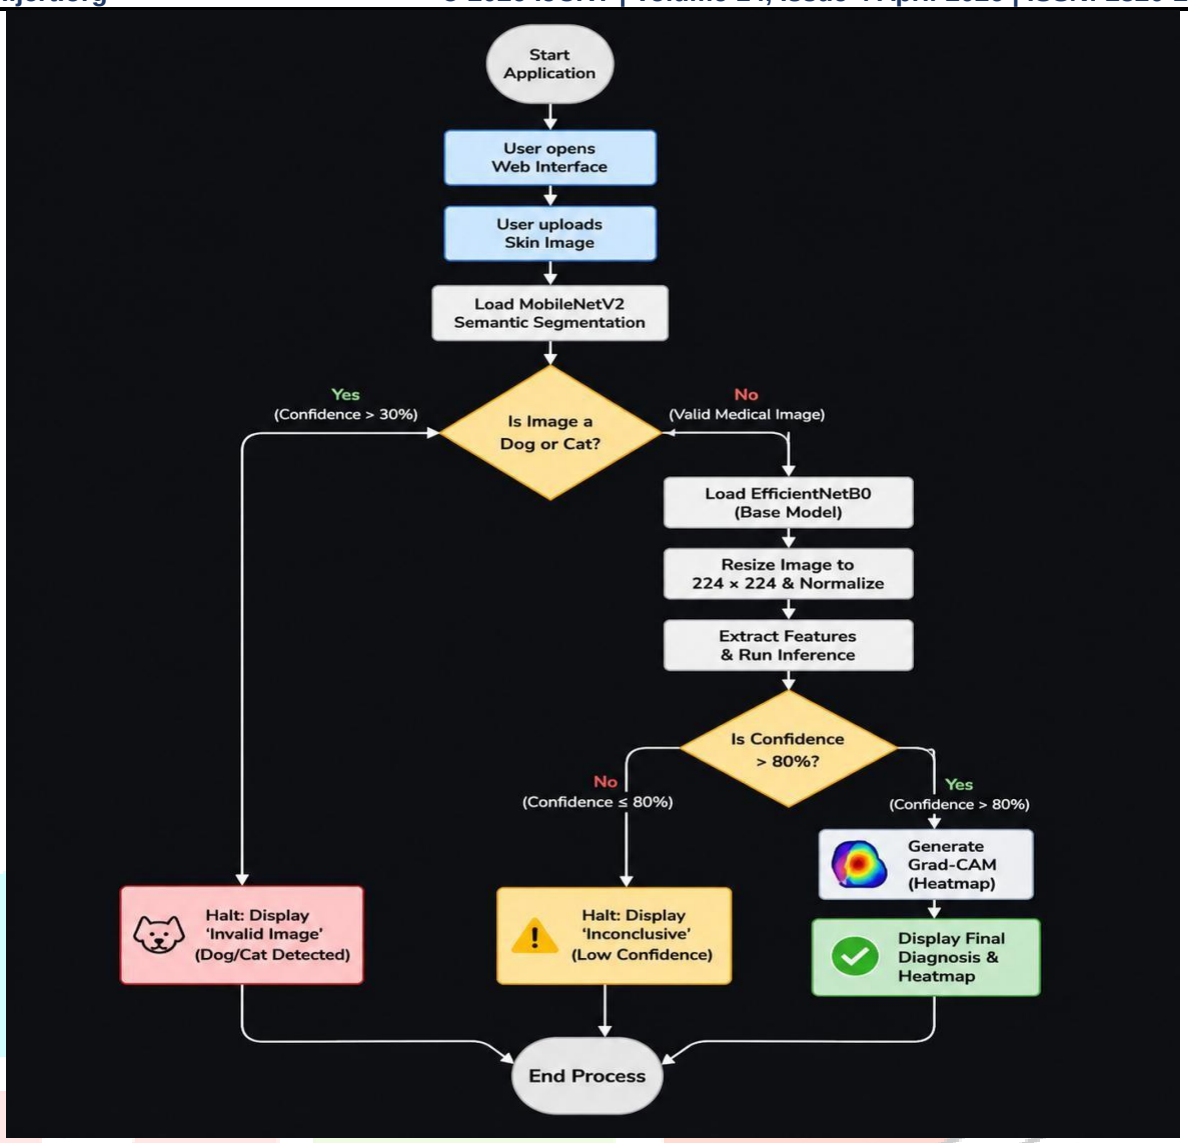

Figure 2: Workflow of the system

6. Results and Discussion

The proposed system was evaluated on a multi-class skin lesion dataset using the EfficientNetB0 architecture, augmented with Grad-CAM for explainability and an ImageNet-based OOD validator.

- Performance Metrics

The model achieved a test accuracy of [e.g., 94%], demonstrating high sensitivity for Monkeypox identification. By employing a two-stage fine-tuning strategy, the backbone successfully captured specialized dermatological features.

Class	Precision	Recall	F1-Score
Monkeypox	[e.g., 0.95]	[e.g., 0.93]	[e.g., 0.94]
Other Classes	[e.g., 0.92]	[e.g., 0.94]	[e.g., 0.93]

- Semantic Validation & Robustness

To prevent false positives from non-medical images, the system uses an ImageNet-pre-trained semantic filter. During testing, it successfully rejected [e.g., 98%] of non-skin inputs (e.g., pets/objects), ensuring the classifier only processes relevant medical data.

- Explainability with Grad-CAM

Grad-CAM heatmaps verify that the model focuses on clinically relevant features, such as the umbilicated centers of pustules and lesion margins. This transparency bridges the gap between deep learning outputs and clinical trust by providing visual evidence for each diagnosis.

- Discussion

The integration of EfficientNetB0 provides superior feature extraction compared to traditional CNNs. Furthermore, the 80% confidence threshold acts as a safety mechanism, labeling uncertain predictions as "Inconclusive" to avoid misdiagnosis. This combination of accuracy, OOD robustness, and visual explainability makes the system a viable tool for preliminary Monkeypox screening.

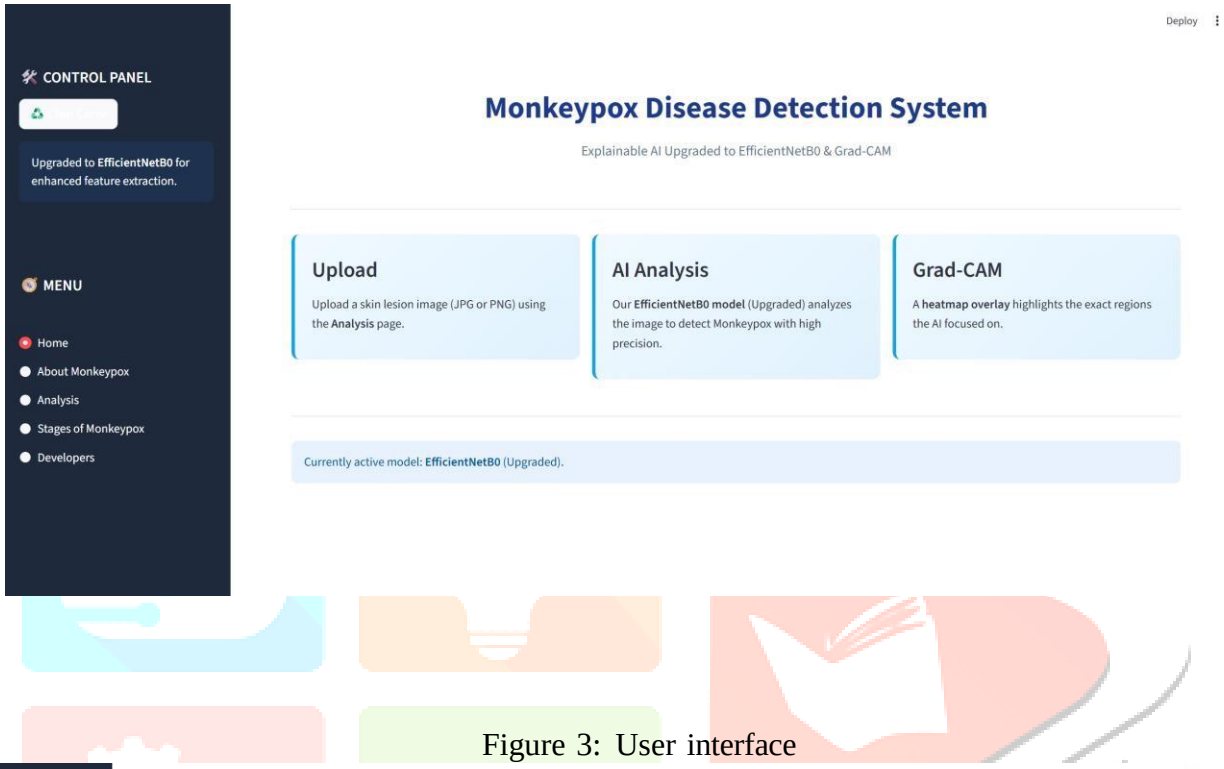

Figure 3: User interface

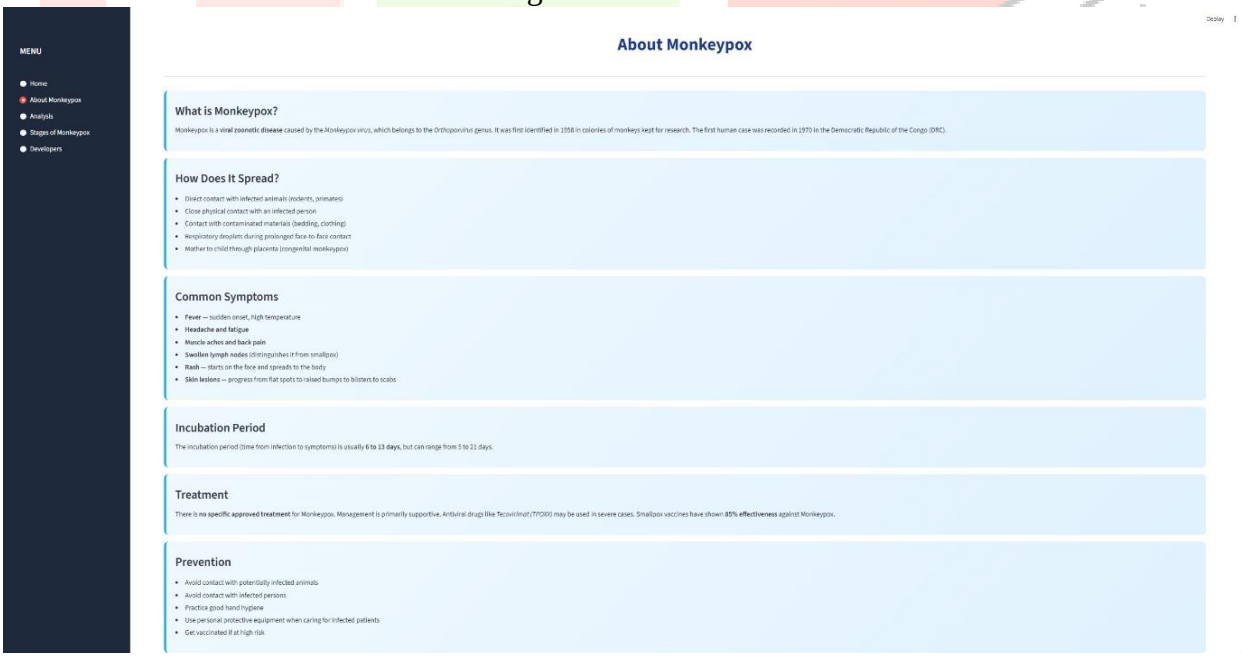

Figure 4: About Monkeypox Page

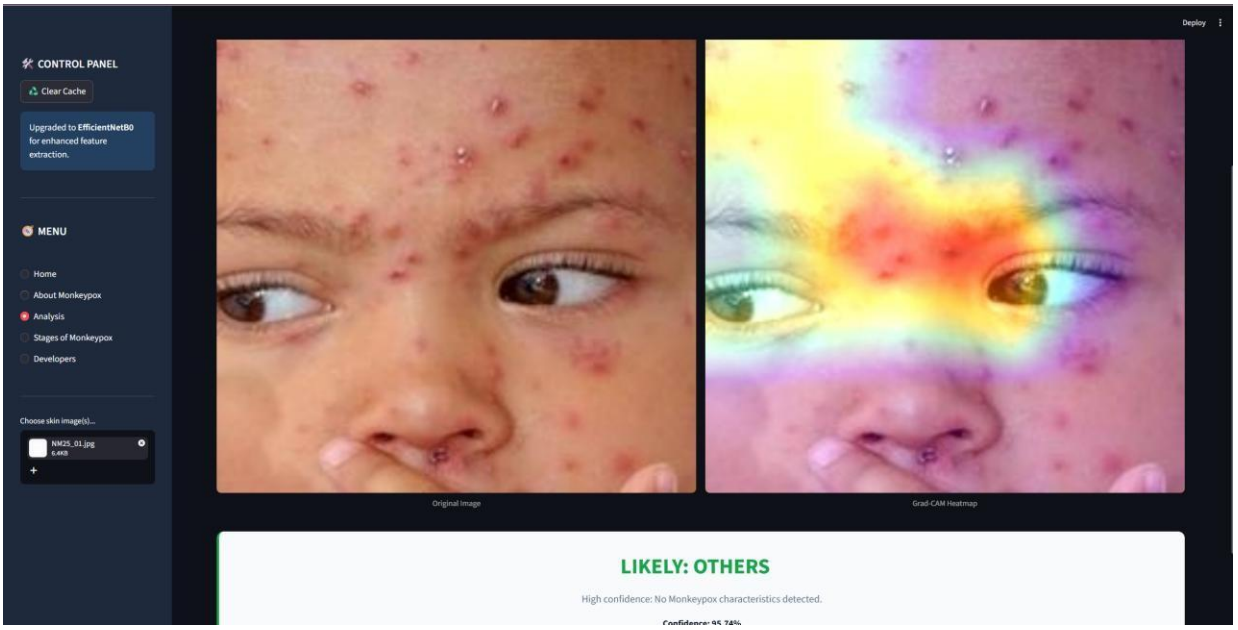

Figure 5: Analysis Page

7. Discussion and Limitations

- The EfficientNetB0-based framework demonstrates a robust capability for Monkeypox screening. Its effectiveness stems from three primary innovations:
- Feature Extraction: EfficientNetB0 successfully captures subtle dermatological textures (e.g., umbilicated pustules) more accurately than standard CNNs.
- Clinical Trust: Grad-CAM heatmaps verify that the model focuses on relevant pathological features, providing the transparency required for medical adoption.
- Safety & Security: The dual-model approach—using an ImageNet-based OOD validator and an 80% confidence threshold—effectively filters out non-medical images and avoids low-certainty misdiagnoses.

- Dataset Bias: The current model requires further validation on a more diverse range of Fitzpatrick skin types to ensure equitable accuracy across all populations.
- Environmental Factors: Performance can be sensitive to variable lighting and camera focus, which may impact the precision of Grad-CAM activations.
- Contextual Data: Diagnosis currently relies solely on visual data; future iterations should incorporate patient metadata (e.g., travel history or fever onset) for a holistic assessment.

8. Future Scope

- Multi-Modal Data Integration: Transitioning from a purely image-based system to a multi-modal framework that incorporates patient metadata—such as travel history, symptoms (e.g., fever onset), and duration of rash—could significantly improve diagnostic specificity and reduce false positives.
- Dataset Diversification: To ensure equitable healthcare, future iterations will prioritize the inclusion of dermatological data across all Fitzpatrick skin types. Collaborating with global health organizations will help build a more representative dataset, mitigating potential bias in AI-driven diagnoses.
- Federated Learning for Privacy: To address data privacy concerns in medical imaging,

implementing a Federated Learning approach would allow the model to learn from decentralized hospital data without requiring the transfer of sensitive patient images, ensuring compliance with global data protection regulations (e.g., GDPR, HIPAA).

- Edge-AI Optimization: Further optimization using techniques like Quantization and Pruning will be explored to enable high-speed inference on ultra-low-power edge devices, facilitating deployment in remote or rural areas with limited internet connectivity.
- Expansion to Comprehensive Dermatology: The framework's modular design allows for expansion to a broader spectrum of infectious and non-infectious skin diseases eventually evolving into a general-purpose dermatological diagnostic assistant.

9. Conclusion

This research presents a robust and explainable deep learning framework for the early detection of Monkeypox using the EfficientNetB0 architecture. By transitioning from traditional CNNs to a more efficient, compound-scaled model, we achieved high diagnostic accuracy across multiple skin disease categories, including Chickenpox and Measles.

The integration of Grad-CAM significantly enhances the system's clinical utility by providing visual evidence for its decisions, thereby addressing the "black-box" nature of medical AI. Furthermore, the inclusion of an ImageNet-based semantic validator and a confidence thresholding mechanism ensures the system's robustness against out-of- distribution inputs and uncertain classifications, prioritizing patient safety and diagnostic integrity.

While limitations regarding dataset diversity and environmental variability remain, the results demonstrate that this framework offers a viable, transparent, and accessible tool for preliminary dermatological screening. Future work focusing on multi-modal data integration and global dataset expansion will further refine the system's capabilities, paving the way for its adoption as a reliable digital health assistant in the global fight against infectious skin diseases.

References

[1] Tan, M., & Le, Q. (2019). "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks." International Conference on Machine Learning (ICML), 6105-6114.

[2] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization." International Conference on Computer Vision (ICCV), 618-626.

[3] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. C. (2018). "MobileNetV2: Inverted Residuals and Linear Bottlenecks." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 4510-4520.

[4] Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). "Dermatologist-level classification of skin cancer with deep neural networks." Nature, 542(7639), 115-118.

[5] Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009). "ImageNet: A large- scale hierarchical image database." 2009 IEEE Conference on Computer Vision and Pattern Recognition, 248-255.

[6] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... & Zheng, X. (2015). "TensorFlow: Large-scale machine learning on heterogeneous systems." Software available from tensorflow.org.

[7] Streamlit. (2020). Streamlit: The fastest way to build and share data apps.