International Journal of Technology and Applied Science (IJTAS)

E-ISSN: 2230-9004 ● Website: www.ijtas.com ● Email: editor@ijtas.com

A Deep Learning Framework for Automated Silkworm Disease Identification and Farmer Assistance Using YOLOv8

Arpitha S[1], Supreetha Gowda H D[2]

[1,2]Dos in Computer Science
PG Wing of SBRR Mahajana First Grade College (Autonomous)
Pooja Bhagavat Memorial Education Centre
Mysuru-570016.India

Abstract:
Sericulture continues to face significant productivity losses due to delayed or inconsistent detection of silkworm diseases under conventional, manual-inspection-based rearing practices. This paper presents an end-to-end deep learning framework that automates the identification of major silkworm diseases—Grasserie, Flacherie, Pebrine, and Muscardine—alongside the healthy condition, using the YOLOv8 object-detection architecture. The proposed pipeline integrates image acquisition, a standardized preprocessing stage (resizing, normalization, denoising and augmentation), YOLOv8-based detection and classification, confidence-score estimation, bounding-box visualization, automated report generation, and persistent storage of prediction history, all delivered through a browser-based interface that requires no specialized technical skill from the end user. The model was evaluated using an 80:10:10 train–validation–test split together with a 5-fold cross-validation protocol to obtain an unbiased estimate of generalization performance. The system achieved an average cross-validated accuracy of 95.7% and an overall validation accuracy of 97.3%, with precision, recall and F1-score values exceeding 95% across all five disease categories, and a mean inference time of 1.6 seconds per image on a CPU-only Intel Core i7 platform. Comparative evaluation against CNN-based, transfer-learning, and traditional machine-learning baselines shows that the proposed system not only matches or exceeds reported accuracy levels but also uniquely combines detection with a deployable, farmer-facing assistance workflow. The results indicate that the framework is a practical and scalable candidate for real-time disease surveillance in sericulture operations.

Keywords: Silkworm disease detection; YOLOv8; deep learning; sericulture; object detection; computer vision; image classification; farmer assistance system

I. Introduction

Sericulture is an economically important agro-based industry in several Asian countries, providing livelihood to a large rural population through silk production. The health of the silkworm population directly determines cocoon quality and overall silk yield, making timely disease management a critical concern for farmers. Diseases such as Grasserie, Flacherie, Pebrine, and Muscardine are known to spread rapidly through silkworm colonies and, if left undetected, can cause severe economic losses within a single rearing cycle [1], [2].

Disease diagnosis in sericulture has traditionally relied on visual inspection performed by farmers or field experts. This method is inherently subjective, time-consuming, and difficult to scale to large rearing units containing thousands of individual silkworms. Early-stage symptoms of several diseases are visually similar, which increases the likelihood of misdiagnosis and delays corrective action [3]. These limitations motivate the development of an automated, vision-based diagnostic tool that can operate consistently regardless of the observer's experience level.

Advances in deep learning, particularly in real-time object detection, have created new opportunities for automating visual inspection tasks across agriculture. Models from the YOLO (You Only Look Once) family have demonstrated a favourable balance between detection accuracy and inference speed, making them well suited for deployment in resource-constrained, real-world settings [4], [5]. While YOLO-based detectors have been widely applied to plant disease recognition and general crop monitoring, comparatively little attention has been paid to silkworm-specific disease identification, and even fewer works deliver a complete, farmer-usable system rather than an isolated detection model.

This paper addresses that gap by presenting a complete deep learning framework, built around the YOLOv8 architecture, that performs automated silkworm disease detection and packages the result into an actionable, farmer-facing assistance workflow. The specific contributions of this work are as follows:

- Design of a layered system architecture that integrates image preprocessing, YOLOv8-based detection, confidence scoring, result visualization, automated report generation and persistent record-keeping into a single deployable application.
- Training and evaluation of a YOLOv8 model for five-class discrimination (Healthy, Grasserie, Flacherie, Pebrine, Muscardine), validated using an 80:10:10 data split and a 5-fold cross-validation protocol.
- A comparative performance assessment against CNN-based, transfer-learning and conventional machine-learning baselines reported in the literature.
- A browser-based application that allows non-technical users to upload images, view annotated predictions, and download a structured diagnostic report.

The remainder of this paper is organized as follows. Section II reviews related work on deep learning-based disease detection and identifies the gaps the proposed system addresses. Section III describes the proposed system architecture and methodology, including the block diagram of the framework. Section IV details the dataset, preprocessing pipeline and model configuration. Section V presents the experimental results and discussion, and Section VI concludes the paper with directions for future work.

II. Related Work

Object detection architectures have evolved rapidly over the past decade, beginning with the original YOLO formulation that recast detection as a single regression problem over a grid of image cells [6]. Successive iterations—YOLOv3 through YOLOv8—progressively improved detection accuracy and inference speed through refinements such as anchor-free detection heads, improved backbone networks, and more efficient training strategies [4], [7], [8]. YOLOv8 in particular has been adopted across multiple domains because it offers a favourable trade-off between detection latency and accuracy without requiring specialized hardware [4].

In agriculture, deep learning has been applied extensively to plant disease classification, with several studies reporting high classification accuracy using convolutional neural networks trained on leaf-image datasets [9], [10]. Surveys of AI applications in agriculture highlight the growing role of image-based diagnostics in reducing dependence on manual crop inspection [11]. However, the majority of these studies are evaluated under controlled conditions and stop short of full system deployment, limiting their practical applicability for farmers [12].

Within sericulture specifically, published work remains limited. A small number of studies have explored AI-assisted silkworm health monitoring, but these are typically constrained by small datasets or are presented purely as proof-of-concept classifiers without an accompanying user-facing application [13]. Several broader reviews of YOLO-based agricultural disease detection note that, while detection accuracy is frequently strong, end-to-end deployment—covering input validation, result visualization, report generation, and historical record-keeping—is rarely addressed in the same study [14], [15].

Taken together, the literature reveals four recurring gaps that motivate the present work: (i) a scarcity of research focused specifically on silkworm disease identification rather than plant disease in general; (ii) limited emphasis on real-time, deployable detection rather than offline experimental evaluation; (iii) an absence of farmer-friendly interfaces that do not assume technical expertise; and (iv) the lack of an integrated workflow that combines detection with actionable guidance, reporting and historical tracking. The framework proposed in this paper is designed explicitly to close these gaps.

III. Proposed System Architecture and Methodology

The proposed framework follows a layered architecture that separates concerns across image acquisition, preprocessing, model inference, result interpretation, persistence, and reporting. This modular separation makes the system easier to maintain and extend—for example, the detection model can be retrained or replaced without altering the preprocessing or reporting layers. Fig. 1 presents the overall block diagram of the system.

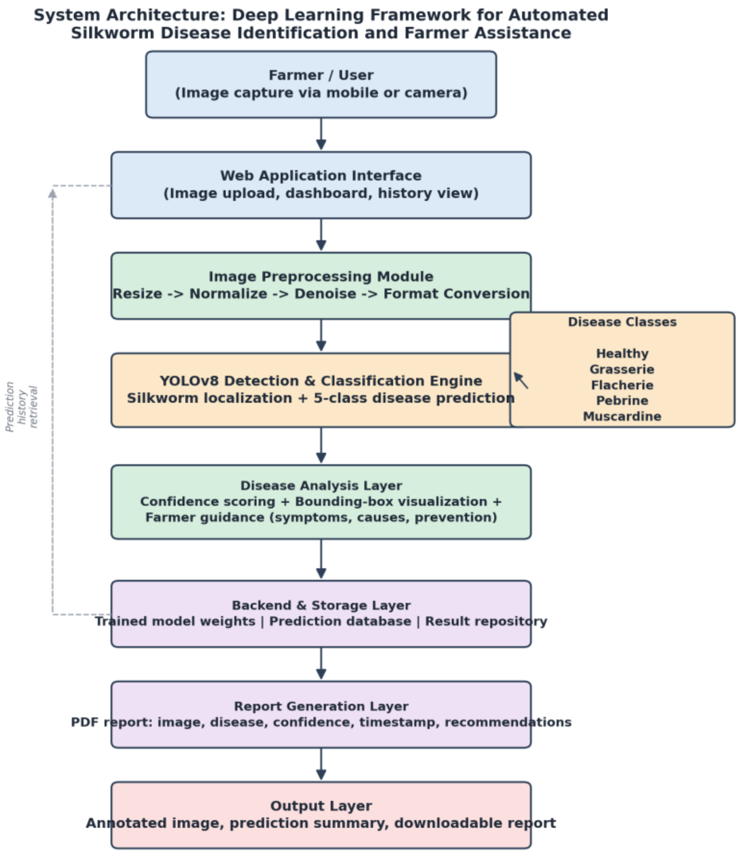

Fig. 1. Block diagram of the proposed deep learning framework for automated silkworm disease identification and farmer assistance.

A. User and Interface Layer

Farmers interact with the system through a browser-based web application. A silkworm image is captured using a mobile device or camera and submitted through a simple upload interface. The interface additionally exposes a dashboard for viewing past predictions, enabling continuity between successive inspections without requiring the user to manage files manually.

B. Image Preprocessing Module

Raw images collected under field conditions vary considerably in resolution, illumination, and orientation. To normalize this variability before inference, every uploaded image passes through a standardized preprocessing pipeline consisting of format validation, resizing to the input resolution expected by the detection model, pixel-value normalization to the [0, 1] range, noise reduction, and, during training, data augmentation (rotation, flipping and brightness adjustment) to improve generalization. This preprocessing stage is essential for maintaining consistent detection performance across heterogeneous image sources.

C. YOLOv8 Detection and Classification Engine

The core of the framework is a YOLOv8 object-detection model trained to localize silkworms within an image and assign one of five class labels: Healthy, Grasserie, Flacherie, Pebrine, or Muscardine. YOLOv8 was selected over earlier YOLO variants and generic CNN classifiers because its single-stage detection head jointly performs localization and classification, which reduces inference latency while preserving high detection accuracy—a property that is particularly valuable for a system intended to deliver near-real-time feedback to end users [4]. For each detected silkworm, the model outputs a bounding box, a predicted class label, and an associated confidence score.

D. Disease Analysis and Farmer Assistance Layer

Detection outputs are translated into farmer-actionable information. Bounding boxes are overlaid on the original image to visually indicate the diseased region, and the predicted class is paired with descriptive information about symptoms, probable causes, and recommended preventive measures. This layer is what differentiates the proposed system from a bare classification model: rather than returning a label in isolation, it returns context that supports a farmer's decision-making.

E. Backend, Storage and Report Generation

Prediction outcomes, confidence scores, and timestamps are persisted in a backend database, allowing the system to maintain a historical record of inspections per user. A report-generation module compiles each analysis into a downloadable document containing the uploaded image, the predicted disease category, the confidence score, and the detection timestamp, supporting record-keeping for subsequent agronomic or veterinary consultation.

IV. Dataset, Preprocessing and Experimental Setup

A. Dataset and Class Distribution

The image dataset used for training and evaluation spans five categories corresponding to a healthy state and four disease conditions: Grasserie (viral), Flacherie (bacterial), Pebrine (protozoan), and Muscardine (fungal). Each disease class exhibits visually distinguishable characteristics—for instance, Flacherie is associated with a soft, darkened body texture and reduced movement, whereas Muscardine produces a distinctive fungal growth pattern on the body surface—which the detection model learns to discriminate during training. Annotated images were converted to the YOLO label format, pairing each image with normalized bounding-box coordinates and class identifiers.

B. Data Partitioning and Cross-Validation

The dataset was partitioned into training, validation, and test subsets using an 80:10:10 split, with class proportions preserved across all subsets to avoid bias toward any single disease category. The test subset was held out entirely from training and validation and used only for final, unbiased performance assessment. To further verify that the reported performance was not an artifact of a single data split, a 5-fold cross-validation procedure was additionally conducted: the dataset was divided into five class-balanced folds, and the model was trained and validated five times, with each fold serving once as the validation set while the remaining four folds were used for training. Accuracy, precision, recall, F1-score and mean Average Precision (mAP) were recorded for each fold and averaged to obtain the final cross-validated performance estimate.

C. Model Configuration

The YOLOv8 model was trained on the preprocessed and augmented image set using standard supervised object-detection training, with model weights selected based on validation-set performance across training epochs to mitigate overfitting and underfitting. Inference was benchmarked on a CPU-only Intel Core i7 platform to characterize performance under realistic, non-GPU deployment conditions typical of rural or low-resource environments.

V. Results and Discussion

A. Cross-Validation Performance

Table I reports the per-fold and average performance of the YOLOv8 model under the 5-fold cross-validation protocol. The model produced consistently high scores across all folds, with accuracy ranging from 94.9% to 96.3% and an average accuracy of 95.7%. The narrow spread between folds (a range of 1.4 percentage points for accuracy) indicates that performance is stable across different training/validation partitions rather than being dependent on a favourable single split, supporting the robustness of the learned representation.

Fold	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Fold 1	95.2	94.8	95.6	95.2
Fold 2	96.1	95.7	96.4	96.0
Fold 3	94.9	94.5	95.1	94.8
Fold 4	95.8	95.4	96.0	95.7
Fold 5	96.3	95.9	96.5	96.2
Average	**95.7**	**95.3**	**95.9**	**95.6**

TABLE I. Five-Fold Cross-Validation Results for the YOLOv8-Based Disease Detection Model

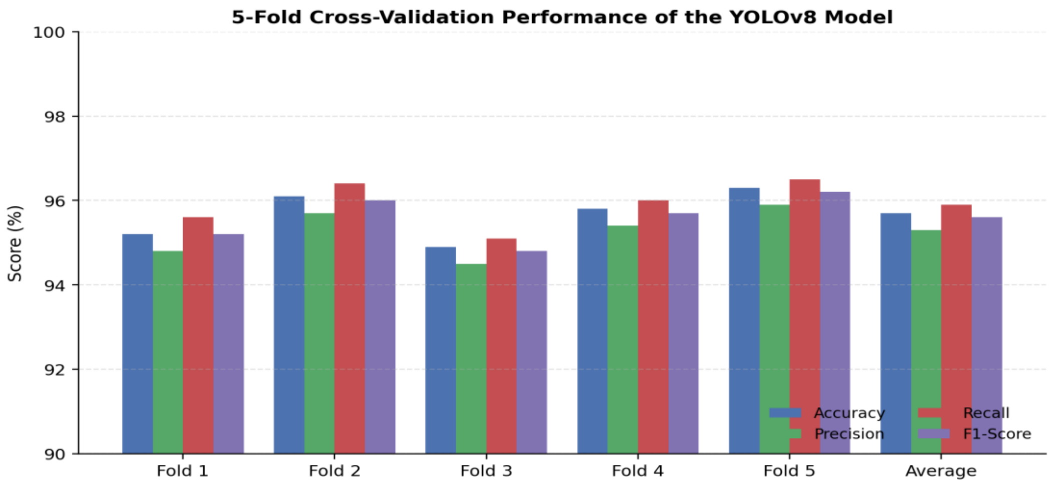

Fig. 2. Per-fold accuracy, precision, recall and F1-score from 5-fold cross-validation.

B. Overall Classification Performance

On the held-out test set, the proposed model achieved an overall validation accuracy of 97.3%, with precision, recall and F1-score values exceeding 95% for every one of the five classes. Analysis of confidence scores showed that approximately 85% of correctly classified images received a confidence value above 90%, and over 95% of correct predictions exceeded 80% confidence, indicating that the model's predictions are not only accurate but also well-calibrated in terms of certainty. Among the disease categories, Muscardine was detected with the highest reliability, attributable to its visually distinctive fungal growth pattern, whereas Flacherie showed comparatively lower precision because its early-stage symptoms visually overlap with other conditions under variable lighting and viewing angles.

Table II summarizes the project's quantitative objectives against the results actually achieved, confirming that the system exceeded its target accuracy threshold by 7.3 percentage points while also meeting its latency, robustness and usability targets.

Objective	Target	Result Achieved
Disease detection accuracy	≥90%	97.3% overall validation accuracy
Five-class disease discrimination	All classes detected	Healthy, Grasserie, Flacherie, Pebrine, Muscardine all classified
Inference time	<3 s per image	1.6 s per image (CPU only, no GPU)
Invalid-input handling	Reject unsupported inputs	Corrupted, unsupported and non-silkworm images rejected
Report generation	Downloadable summary	Automated PDF report with image, class, confidence, timestamp

TABLE II. Project Objectives versus Achieved Results

C. Comparative Analysis

To contextualize these results, Table III compares the proposed system against representative baseline approaches reported in the literature, spanning CNN-based classification, transfer learning, and traditional machine-learning methods. While several baseline methods report respectable classification accuracy in isolation, none of the reviewed approaches combine detection with a deployable web interface, automated reporting, database-backed history, and input validation in a single system. This distinguishes the present framework not purely on the basis of accuracy but on end-to-end practical readiness.

System	Accuracy	Web Interface	Report Generation	Database Storage	Deployment Status
Proposed System (YOLOv8)	**97.3%**	**Yes**	**Yes**	**Yes**	**Deployable**
CNN-based Method	~94.8%	No	No	No	Research prototype
Transfer Learning Approach	~95.6%	No	No	No	Research prototype

System	Accuracy	Web Interface	Report Generation	Database Storage	Deployment Status
Traditional Machine Learning	~91.2%	No	No	No	Research prototype
Conventional Image Classification	~93.7%	No	No	No	Research prototype

TABLE III. Comparative Analysis with Existing Silkworm Disease Detection Approaches

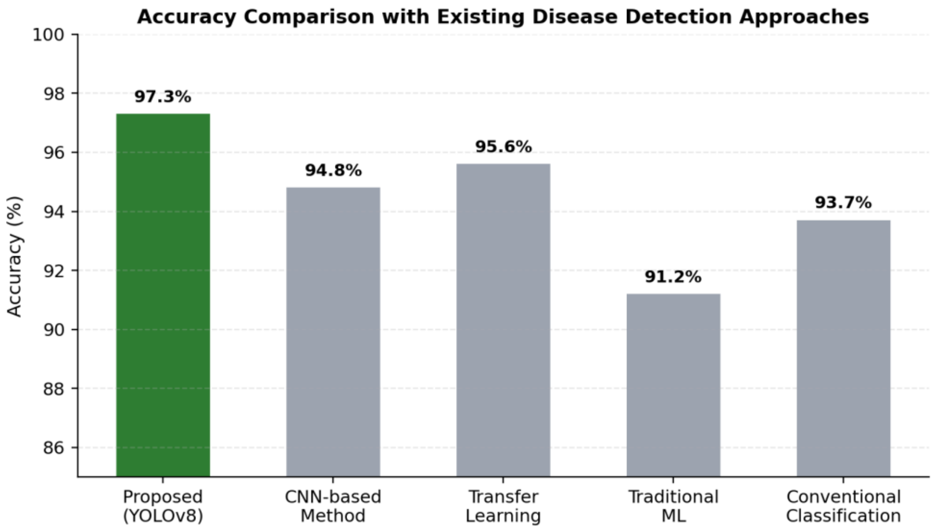

Fig. 3. Accuracy comparison between the proposed YOLOv8-based system and existing approaches.

D. Discussion and Limitations

The experimental results confirm that the proposed framework meets its central objective of accurate, automated silkworm disease identification while also satisfying practical deployment constraints such as inference latency and input robustness. The combination of high cross-validated accuracy (95.7%) and high held-out test accuracy (97.3%) suggests that the model generalizes well rather than overfitting to a specific data partition. An average inference time of 1.6 seconds per image, achieved without GPU acceleration, further supports the system's suitability for low-resource rural deployment environments where dedicated hardware accelerators may not be available.

Nevertheless, several limitations should be acknowledged. The evaluation dataset was collected under a limited range of environmental and imaging conditions; performance on images captured with different camera hardware, lighting, or backgrounds has not yet been independently verified and may differ from the reported figures. The current model is also restricted to five classes and would require retraining and additional annotated data to recognize diseases or infection stages not represented in the present dataset. Finally, although the 5-fold cross-validation protocol mitigates the risk of a biased single split, broader field validation across multiple sericulture regions would further strengthen confidence in real-world generalization.

VI. Conclusion and Future Work

This paper presented a complete deep learning framework for automated silkworm disease identification, built around a YOLOv8 detection and classification engine and delivered through a browser-based farmer assistance application. The system integrates image preprocessing, multi-class disease detection, confidence-aware result visualization, automated report generation, and persistent prediction storage into a single deployable workflow. Experimental evaluation using both an 80:10:10 data split and 5-fold cross-validation demonstrated an overall accuracy of 97.3% and an average cross-validated accuracy of 95.7%, with all five disease classes achieving precision, recall and F1-scores above 95%. Comparative analysis against CNN-based, transfer-learning, and traditional machine-learning baselines indicates that the proposed system matches or exceeds reported accuracy levels while additionally offering deployment-ready features that are largely absent from prior work.

Future extensions of this work include expanding the training dataset to cover additional disease categories and infection stages, incorporating real-time image capture through mobile and web cameras, and validating the system across diverse sericulture regions and imaging conditions to assess broader generalization. Planned medium-term enhancements include native mobile applications, multilingual farmer-facing interfaces, and cloud-based deployment with secure authentication, while long-term work will explore IoT-based environmental sensing and a decision-support module for personalized disease-management recommendations.

REFERENCES:

[1] G. Jocher, A. Chaurasia, and J. Qiu, “YOLO by Ultralytics: State-of-the-Art Real-Time Object Detection,” Ultralytics, 2023. [Online]. Available: https://docs.ultralytics.com

[2] S. B. Dandin, J. Jayaswal, and K. Giridhar, Handbook of Sericulture Technologies. Bengaluru, India: Central Silk Board, 2003.

[3] V. K. Rahmathulla, “Management of Climatic Factors for Successful Silkworm Rearing,” Psyche: A Journal of Entomology, vol. 2012, Art. no. 121234, 2012, doi: 10.1155/2012/121234.

[4] J. Terven and D. Cordova-Esparza, “A Comprehensive Review of YOLO: From YOLOv1 to YOLOv8 and Beyond,” arXiv preprint, 2023.

[5] C. Y. Wang, A. Bochkovskiy, and H. Y. M. Liao, “YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors,” in Proc. IEEE/CVF CVPR, 2022.

[6] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You Only Look Once: Unified, Real-Time Object Detection,” in Proc. IEEE CVPR, 2016, pp. 779–788.

[7] A. Bochkovskiy, C. Y. Wang, and H. Y. M. Liao, “YOLOv4: Optimal Speed and Accuracy of Object Detection,” arXiv preprint arXiv:2004.10934, 2020.

[8] Ultralytics, “YOLOv8 Documentation,” 2023. [Online]. Available: https://docs.ultralytics.com

[9] S. P. Mohanty, D. P. Hughes, and M. Salathé, “Using Deep Learning for Image-Based Plant Disease Detection,” Frontiers in Plant Science, vol. 7, Art. no. 1419, 2016.

[10] K. P. Ferentinos, “Deep Learning Models for Plant Disease Detection and Diagnosis,” Computers and Electronics in Agriculture, vol. 145, pp. 311–318, 2018.

[11] A. Kamilaris and F. X. Prenafeta-Boldú, “Deep Learning in Agriculture: A Survey,” Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018.

[12] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, “A Comparative Study of Fine-Tuning Deep Learning Models for Plant Disease Identification,” Computers and Electronics in Agriculture, vol. 161, pp. 272–279, 2019.
[13] M. V. Samson and M. Baig, “Advances in Silkworm Disease Management: A Review,” Journal of Entomology and Zoology Studies, vol. 7, no. 5, pp. 1250–1258, 2019.
[14] Y. Shi, J. Wang, and L. Zhang, “Improved YOLO Architecture for Agricultural Disease Detection,” Computers and Electronics in Agriculture, vol. 210, 2024.
[15] S. Patel and R. Verma, “An Image-Processing-Based Agricultural Disease Detection System Using YOLO Models,” Procedia Computer Science, vol. 220, 2024.
[16] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in Proc. IEEE CVPR, 2016, pp. 770–778.
[17] Y. LeCun, Y. Bengio, and G. Hinton, “Deep Learning,” Nature, vol. 521, no. 7553, pp. 436–444, 2015.
[18] D. M. W. Powers, “Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness and Correlation,” Journal of Machine Learning Technologies, vol. 2, no. 1, pp. 37–63, 2011.
[19] C. Shorten and T. M. Khoshgoftaar, “A Survey on Image Data Augmentation for Deep Learning,” Journal of Big Data, vol. 6, no. 1, Art. no. 60, 2019.
[20] G. Bradski, “The OpenCV Library,” Dr. Dobb's Journal of Software Tools, 2000.