



International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

EfficientNet-Based Deep Learning for Medicinal Plant Identification and Leaf Disease Classification: A Comprehensive Review

Anjali Shrishail Bhalkhede¹, Ashlesha Sandeep Mandhare², Gauri Subhash Salunkhe³, Prof. Vaishali Yeole⁴

¹⁻³ Student, ⁴Guide

Department of Computer Engineering, Smt. Indira Gandhi College of Engineering, Navi Mumbai, India

ABSTRACT—

Distinguishing one medicinal plant from another by its leaves and spotting diseases on those leaves are two related but separate computer vision problems. EfficientNet has been used as a backbone in both studies. In this review, we examine how EfficientNet and the B0 baseline have been used for medicinal plant identification and leaf disease classification and set that work against the wider background of convolutional architectures, transfer learning, datasets, and explainability. We compiled a pool of 60 candidate sources covering foundational architecture and transfer-learning papers, plant disease studies, medicinal plant studies, datasets, explainability methods, and earlier reviews. Eleven plant-focused papers were examined against six research questions, and the rest served as background information until they could be verified. The studies we examined reported very high accuracy for EfficientNet variants, up to 99.97% in one disease study and 99.96% in one medicinal plant study. Because the variants, datasets, split protocols, and metrics all differ, these figures cannot be ranked against one another. However, several concerns remain: heavy reliance on controlled benchmark data, small medicinal plant datasets, limited reporting on explainability, and almost no work that combines species identification with leaf-health assessment. We discuss these gaps and suggest directions for an integrated system, such as BotaniQ.

Keywords—EfficientNet, EfficientNet-B0, medicinal plant identification, plant disease classification, transfer learning, Grad-CAM, and literature review.

1. INTRODUCTION

Medicinal plants still play a significant role in healthcare, and leaves are among the plant parts most often used in preparations [6]. Getting the identification right matters for safe use, but it is harder than it sounds. Many species have similar-looking foliage, and it usually requires specialist knowledge to distinguish them [5], [8]. Leaf disease diagnosis is also difficult. Most diseases leave visible symptoms; however, diagnosis has traditionally relied on an experienced pathologist looking at the leaf, which is slow and depends on the person who is looking [1].

Convolutional neural networks (CNNs) have been applied to both problems on a large scale, mainly because they learn visual features directly from images instead of depending on hand-designed descriptors [1], [21]. Of the available architectures, EfficientNet [12] has garnered considerable attention. It scales network depth, width, and input resolution together, so that accuracy can be traded against computational cost in a controlled manner. Its smallest member, EfficientNet-B0, is especially attractive when the model must run on modest hardware, a motivation that at least one herbal-leaf study states outright [8].

Two decisions shape the scope of this review. First, we treat medicinal plant identification and plant disease classification as separate tasks. The labels mean different things (a species in one case and a health condition in the other), the visual variation the model has to cope with is different, and so is the cost of getting an answer wrong.

Second, we do not take high reported accuracy as proof that a model is the best choice. Studies differ in their datasets, splits, and metrics, and several rely on controlled image collections that may not look much like field conditions [3], [21], [22].

The review is organized around six research questions (Section III-A). It offers a structured comparison of EfficientNet-based studies for both tasks, a critical look at what the reported results do and do not show, and a set of research gaps that matter for integrated identification and leaf health systems such as BotaniQ. Section II provides the literature background, Section III explains how the review was conducted, Section IV presents the results and discussion, and Section V concludes the paper.

2. LITERATURE REVIEW

A. Medicinal plant identification and leaf disease classification

Identifying a medicinal plant is a species-level recognition problem. Recent work has pointed to the same obstacle: leaves from different species look alike, and there are few comprehensive image databases to train on [5]. The datasets that do exist tend to be modest; for example, 3,500 images across ten herbal leaf classes [6] or 500 manually collected images [8]. Some studies have gone beyond species labels and attempted to identify plant parts and traditional uses as well [4].

Leaf disease classification asks a different question: given a plant, which health condition is present? Models must separate subtle symptom patterns within or across crops, and the literature notes that fine-grained symptom differences, limited generalizability, and weak interpretability are still open problems [3]. The two tasks have different label spaces and different risks; therefore, the results from one cannot be simply carried over to the other. Broader surveys of deep learning in agriculture [28] and species identification by computer vision [31] provide useful context for both.

B. CNN architectures and transfer learning

Plant-image studies tend to compare the same small set of established architectures: VGG [15], ResNet [14], DenseNet [16], Inception, and MobileNet variants [17]. Training these from scratch usually requires more data than plant datasets can offer; therefore, most studies begin with ImageNet-pretrained weights. Transfer learning has a broad theoretical treatment [18], and empirical work shows that the extent to which learned features transfer depends on whether a layer is general or task-specific [19]. Later work asked whether architectures that perform better on ImageNet also transfer better to other tasks [20]. In the plant domain, early deep-learning studies reached high disease classification accuracy on a controlled public dataset [21], while later work stressed that dataset size and variety strongly affect how well such models perform [22].

C. EfficientNet and compound scaling

EfficientNet was introduced as a family of networks built by scaling a baseline model, EfficientNet-B0, with a compound coefficient that increases the depth, width, and resolution [12]. The baseline uses inverted bottleneck blocks popularized by MobileNetV2 [17]. A later family, EfficientNetV2, was proposed to train faster and obtain smaller models [13]. These details matter here because the reviewed papers use different members of the family (B0, B3, B4/B5, V2-B0, and V2-B1), and a result from one variant does not automatically hold for another.

D. Datasets, explainability and confidence

PlantVillage is the public dataset most often cited for disease classification [25], and the benchmark results on it are frequently very high. Field-condition datasets, such as PlantDoc [30], were created partly because laboratory-style images do not reflect real use. From the interpretability perspective, gradient-based localization methods, such as Grad-CAM [24], highlight the image regions that drive a prediction, and symptom-visualisation work on tomato diseases applied the same idea to plant images [23]. Separately, modern neural networks can produce poorly calibrated confidence scores [26], which is important for any system that shows probabilities or top-k predictions to users.

3. METHODOLOGY (REVIEW OF LITERATURE)

A. Research questions

The review addresses six questions. RQ1: Which deep learning architectures have been used for medicinal plant identification? RQ2: How is transfer learning applied to plant and medicinal plant image classification? RQ3: How has EfficientNet, especially EfficientNet-B0, been evaluated for plant species identification and leaf disease classification? RQ4: Which datasets, preprocessing techniques, augmentation strategies, and evaluation metrics are used? RQ5: What limits reliability, generalization, interpretability, and deployment? RQ6: What research opportunities remain for integrated medicinal plant identification and leaf health assessment?

B. Source identification

We found candidate sources in two ways. First, we ran keyword searches on October 5, 2026 with a general-purpose web search engine, using queries such as “EfficientNet plant disease classification” and “medicinal plant leaf identification deep learning EfficientNet.” Second, we added foundational papers on architectures, transfer learning, datasets, and explainability, along with earlier reviews, from our own background knowledge. This gave us a pool of 60 candidate sources. [TO COMPLETE: list the bibliographic databases searched (for example IEEE Xplore, Scopus, Web of Science, ScienceDirect, Google Scholar), the full search strings and the publication period.]

C. Inclusion and exclusion criteria

For the plant-focused sources, we included studies that applied deep learning to leaf or plant images for species or medicinal plant identification, or for disease classification, and that reported at least a model description and some evaluation outcome. EfficientNet-based studies came first. Foundational papers were retained as background when they defined a method, dataset, or architecture that plant studies relied on. We excluded sources that did not involve images, sources for which no accessible description could be found, and duplicates. An undergraduate thesis version of a journal article [6], for instance, was not counted separately.

D. Screening and data extraction

Screening was informal, and no PRISMA-style counts were maintained; therefore, none are reported here. [TO COMPLETE if a formal screening is carried out.] For each source, a spreadsheet records the title, authors, year, venue, DOI or link, category, priority, whether EfficientNet, medicinal plants, disease classification, transfer learning, or explainability are involved, the dataset, the model, the reported accuracy, and the reason for the source. Anything not visible in the available abstract or record was marked as “Not extracted” and was not estimated.

E. Quality appraisal and scope of the analysis

We appraised the studies qualitatively using five questions. Is the dataset described? Is the split or validation protocol reported? Are more than one metric given? Are baseline models compared? Are limitations discussed? To date, 11 plant-focused sources [1]-[11] have been examined through their abstracts and repository records. Full-text extraction of dataset sizes, splits, baselines, and author-stated limitations is not finished; therefore, the appraisal is limited to what the abstracts say. The remaining sources [12]-[32] are used as background and still need bibliographic verification.

4. RESULTS AND DISCUSSION

A. Composition of the reviewed pool

Of the 60 candidate sources, nine are primary EfficientNet studies on plants: three on disease classification [1]-[3] and six on medicinal or herbal plant identification [4]-[9]. Two additional reviews have been published, one on medicinal plant AI [10] and the other on leaf disease classification [11]. The remainder were architecture papers, transfer learning studies, datasets, explainability methods, and general surveys. This split suggests that direct EfficientNet evidence for medicinal plants is recent and comes mostly from regional or smaller venues, whereas the disease literature leans more toward earlier CNN studies. Table I summarizes these nine studies.

B. RQ1: Architectures used for medicinal plant identification

In the medicinal plant studies examined, EfficientNet variants dominate: V2-B1 [5], V2-B0 [6], [7], B0 [8], an ensemble [4], and one unspecified EfficientNet [9]. Part of that is simply how the pool was built. EfficientNet was a search term, so this should not be read as a sign that other architectures are rarely used. Reviews of medicinal plant AI describe a wider range of approaches, including CNN pipelines and feature-enhancement methods [10], and DeepHerb, a system built on Xception features, is a commonly cited alternative [27]. A balanced answer to RQ1 would require a search of medicinal plant studies that does not restrict the architecture, and that remains outstanding.

C. RQ2: Transfer learning

Transfer learning was the default approach in the studies examined. Atila et al. trained EfficientNet and competing models with transfer learning [1], Adnan et al. tested pretrained CNNs against an EfficientNetB3-based approach [2], and the model of Ashraf et al. is described as a transfer-learning model on an EfficientNetV2 backbone [5]. The herbal-leaf B0 study describes fine-tuning on a limited dataset [8]. None of the abstracts we looked at say how many layers were frozen or what learning-rate schedule was used. These choices matter, because feature transferability changes with layer depth [19], so they should be recorded from the full texts during extraction.

D. RQ3: Evaluation of EfficientNet

For disease classification, the clearest comparative evidence in the examined set is the study by Atila et al., where the larger variants B4 and B5 gave the best results among the models compared [1]. This finding pertains to the larger members of the family, not B0, and the abstract does not mention where B0 ended up. For medicinal plants, B0 appears directly in only one examined study, trained on 500 images [8]. The V2-B0 [6], [7] and V2-B1

[5] studies used a different generation of the architecture [13]. Therefore, evidence specific to EfficientNet-B0 on medicinal plants is thin, and any conclusion about it should be treated as a hypothesis for BotaniQ to test, not an established result.

The reported figures are high across the board: 99.97% and 99.91% accuracy for disease classification [1], 99.96% for medicinal species classification [4], above 97% validation accuracy for herbal leaves [7], and F1 stated as approximately one for Bangladeshi medicinal plants [5]. Numbers this close to the ceiling call for caution. They come from different datasets, class counts, and protocols, and in several cases, the abstracts do not say whether the test data were fully separate from the training data. Differences between these figures should not be read as differences in model quality. Some secondary sources go further and say that EfficientNet outperforms all other CNN models for leaf disease classification [11]. That is a stronger claim than the primary evidence we looked at can support, since the underlying comparison came from a single study design [1].

E. RQ4: Datasets, preprocessing, augmentation and metrics

Dataset sizes in medicinal plant studies range from small to moderate: 3,500 images in ten classes [6], 500 images [8], and several hundred images per class in another study [9]. Disease studies draw on larger benchmark sources, with a 52-category open dataset used in one case [2]. PlantVillage is still the main public reference [25], and the tomato disease study raises concerns about leaning on controlled benchmarks [3]. Preprocessing and augmentation are described in very uneven detail: resizing to 224×224 with rotation and flipping [8], adaptive augmentation [2], and separate evaluation of original and augmented datasets [1]. The last design deserves a closer look. If augmented copies of one image end up in both the training and test partitions, the results can be inflated by leakage, and the abstract does not say how this was avoided. A general survey of augmentation methods [29] is a handy reference for describing these choices.

Metrics vary too. Most studies report accuracy, one adds precision [1], another gives macro and weighted F1 [5], and a third reports precision, recall, and F1 together [9]. Accuracy alone can hide poor performance on minority classes; therefore, studies on imbalanced medicinal plant collections should report per-class or macro-averaged measures. Computational efficiency is the reason people choose EfficientNet, yet none of the abstracts quantified it, although one study reported reduced computational resources as an outcome [9].

F. RQ5: Reliability, generalisation, interpretability and deployment

Three limitations came up repeatedly. The first is generalization. Controlled datasets can produce very high accuracy that may not carry over to field images [3], [21], and dataset size and variety strongly influence the results [22]. Field-condition datasets, such as PlantDoc, are built for this reason [30]. The second is interpretability. The abstracts of the EfficientNet studies we examined [1]-[9] report no explainability analysis, even though methods such as Grad-CAM are well established [24] and have already been applied to plant disease images [23]. We would need to check the full texts before concluding that these studies leave it out. The third is deployment. Mobile use is cited as a motivation [8], and one study reports a public application [7], but we found no systematic measurement of latency, model size, or on-device accuracy in the abstracts. Confidence calibration matters for any system that shows confidence scores or top-k predictions, and none of the examined studies discussed it, although general work shows that modern networks can be poorly calibrated [26].

G. RQ6: Research opportunities

None of the examined studies [1]-[9] combines species identification with leaf-health assessment in one pipeline, although we have not confirmed this from the full texts. Disease studies are organized by crop and symptom, medicinal plant studies by species; therefore, a combined system faces an open design question: should it use one multi-task model, or a species stage followed by species-specific health classification? Table II connects the main gaps to their evidence and to future work.

Transformer-based and hybrid models are another direction, since vision transformers [32] are now used alongside CNNs. However, the pool holds too few verified plant studies of this kind to draw any conclusions.

H. Limitations of this review

This review has several limitations. Only 11 sources were examined at the abstract level, so the extracted values are incomplete, and the appraisal of methodological quality is provisional. The search was not a formal multi-database systematic search, and no screening counts were recorded. The candidate pool was built around EfficientNet, which biases the findings toward that architecture. The background sources [12]-[32] were cited from general knowledge and still have to be verified. Finally, the reported accuracies come from very different experiments, and we place them side by side only to describe them, not to rank them.

TABLE I. EfficientNet-based plant studies examined (values as reported in abstracts or records; “Not extracted” indicates that the record has not yet been read).

Ref.	Task	EfficientNet variant	Data as reported	Reported result	Remarks
[1]	Disease classification	B0-B7 compared; B4 and B5 reported highest	Original and augmented versions of a leaf-image dataset (size not extracted)	B5 / B4 (original / augmented): accuracy 99.91% / 99.97%; precision 98.42% / 99.39%	Compared with other deep models using transfer learning
[2]	Disease classification	B3 with adaptive augmentation	Open dataset, 52 disease and healthy categories	Not extracted	Compared Xception, InceptionResNetV2, InceptionV3, ResNet50
[3]	Disease classification (tomato)	EfficientNet with CBAM and dynamic prototype memory	Not extracted	Not extracted	Notes reliance on controlled benchmark datasets as a concern
[4]	Medicinal species, parts and uses	EfficientNet-based majority-vote ensemble	Ethiopian medicinal plants (size not extracted)	99.96% test and validation accuracy	Ensemble and transfer learning
[5]	Medicinal plant recognition	EfficientNetV2_B1	Three Bangladeshi medicinal plant datasets	Macro and weighted F1 stated as approximately one	Customised transfer-learning model
[6]	Herbal leaf classification	EfficientNetV2B0	10 classes, 350 images each (3,500)	Not extracted	Journal article; thesis version not counted
[7]	Herbal plant classification	EfficientNetV2B0	Augmented foliage images (details not extracted)	Validation accuracy above 97%	Includes a public identification application

Ref.	Task	EfficientNet variant	Data as reported	Reported result	Remarks
[8]	Herbal leaf classification	EfficientNet-B0 (fine-tuned)	500 manually collected images, 224x224, rotation and flipping	Not extracted	Motivated by mobile or resource-limited deployment
[9]	Medicinal plant detection	EfficientNet (variant not extracted)	Several hundred images per class	Not reported in abstract; precision, recall, F1 used	Repository entry; peer-review status unverified

TABLE II. Research gaps, supporting evidence, and suggested directions

Gap	Evidence in reviewed sources	Suggested direction	Relevance to BotaniQ
Limited evidence for EfficientNet-B0 on medicinal plants	One B0 herbal-leaf study with 500 images [8]; larger variants best in a disease study [1]	Controlled comparison of B0 against ResNet, DenseNet and MobileNet on the same data and split	Justifies testing B0 instead of assuming it
Controlled versus field images	Benchmark reliance noted [3]; dataset variety matters [22]; field datasets exist [30]	External validation on field-condition images	Required for real user photos
Evaluation inconsistency	Near-ceiling accuracy under different protocols [1], [4], [5], [7]	Report splits, leakage controls and macro metrics	Supports credible reporting of BotaniQ results
Sparse explainability reporting	No explainability reported in the abstracts [1]-[9]; Grad-CAM available [24]	Systematic Grad-CAM analysis, including failure cases	Planned feature
Confidence and top-k reliability	Calibration issues in modern networks [26]; not discussed in examined plant studies	Evaluate calibration and top-k accuracy	Planned confidence scores and top-k output
No integrated identification and health pipeline	None found in examined set [1]-[9] (full-text check pending)	Compare multi-task and two-stage designs	Core aim of the project
Deployment evidence	Mobile motivation [8]; application reported [7]; latency not extracted	Measure model size, latency and on-device accuracy	Relevant to B0 choice

5. CONCLUSION

This review examined the use of EfficientNet for medicinal plant identification and leaf disease classification, keeping the two as distinct tasks. The studies we examined reported very high accuracy, but the variants, datasets, protocols, and metrics differed so much that the results did not show which model was best. Direct evidence for EfficientNet-B0 on medicinal plants is limited to small datasets, explainability and calibration were rarely reported in the abstracts we read, and we found no integrated identification and health-assessment system among the studies examined. For BotaniQ, EfficientNet-B0 is a reasonable candidate because it is efficient. Its suitability should still be tested against comparable baselines on a well-documented dataset, with external validation, Grad-CAM analysis, and calibration checks. The next step before submission is to finish full-text extraction and reference verification for all 60 sources.

REFERENCES

- [1] Ü. Atila, M. Uçar, K. Akyol, and E. Uçar, "Plant leaf disease classification using EfficientNet deep learning model," *Ecol. Informat.*, vol. 61, Art. no. 101182, Mar. 2021, doi: 10.1016/j.ecoinf.2020.101182.
- [2] F. Adnan, M. J. Awan, A. Mahmoud, H. Nobanee, A. Yasin, and A. M. Zain, "EfficientNetB3-adaptive augmented deep learning (AADL) for multi-class plant disease classification," *IEEE Access*, vol. 11, pp. 85426–85440, 2023, doi: 10.1109/ACCESS.2023.3303131.
- [3] E. Jansi and B. R. Kavitha, "EfficientNet-CBAM-prototype: An attention-guided EfficientNet framework with dynamic prototype representation for tomato disease classification," *Front. Plant Sci.*, vol. 17, Art. no. 1918175, 2026, doi: 10.3389/fpls.2026.1918175.
- [4] M. A. Kiflie, D. P. Sharma, M. A. Haile, and R. Srinivasagan, "EfficientNet ensemble learning: Identifying Ethiopian medicinal plant species and traditional uses by integrating modern technology with ethnobotanical wisdom," *Computers*, vol. 13, no. 2, Art. no. 38, 2024, doi: 10.3390/computers13020038.

- [5] A. B. Ashraf, R. H. Rabbi, B. Biswas, and S. M. T. Siddiquee, "A customized robust deep learning approach for efficient medicinal plant recognition," in *Proc. Int. Conf. Intell. Data Anal. Appl. (IDAA 2025)*, Dhaka, Bangladesh, Dec. 2025. Atlantis Press. [Online]. Available: <https://www.atlantispress.com/proceedings/idaa-25/126024968>
- [6] R. P. S. Putra, C. S. K. Aditya, and G. W. Wicaksono, "Herbal leaf classification using deep learning model EfficientNetV2B0," *JITK (Jurnal Ilmu Pengetahuan dan Teknologi Komputer)*, vol. 9, no. 2, pp. 301–307, 2024, doi: 10.33480/jitk.v9i2.5119.
- [7] A. Sonita *et al.*, "Herbal plant classification using EfficientNetV2B0 model and CRISP-DM approach," *J. Tek. Inform. (JUTIF)*, vol. 6, no. 4, pp. 2253–2262, Aug. 2025, doi: 10.52436/1.jutif.2025.6.4.5141.
- [8] A. N. A. Alda, D. Indra, and F. Umar, "Classification of herbal leaves using EfficientNetB0," *Jurnal CoreIT*, vol. 11, no. 2, pp. 223–232, Dec. 2025, doi: 10.24014/coreit.v11i2.38017.
- [9] P. Sridhar, "EfficientNet for enhanced detection of medicinal plants: A deep learning approach," in *Smart System for Integrated Computing and Communication*. Singapore: Springer, 2026, pp. 177–190.
- [10] M. A. Kiflie, D. P. Sharma, and M. A. Haile, "Deep learning for medicinal plant species classification and recognition: A systematic review," *Front. Plant Sci.*, vol. 14, Art. no. 1286088, 2024, doi: 10.3389/fpls.2023.1286088.
- [11] Z. Y. Younes and E. M. T. A. Alsaadi, "A review on plant leaf disease classification using deep learning," *Iraqi J. Sci.*, vol. 65, no. 11, pp. 6738–6752, Nov. 2024, doi: 10.24996/ij.s.2024.65.11.43.
- [12] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. 36th Int. Conf. Mach. Learn. (ICML)*, 2019, pp. 6105–6114.
- [13] M. Tan and Q. V. Le, "EfficientNetV2: Smaller models and faster training," in *Proc. 38th Int. Conf. Mach. Learn. (ICML)*, 2021, pp. 10096–10106.
- [14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778.
- [15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015.
- [16] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 4700–4708.
- [17] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018, pp. 4510–4520.
- [18] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [19] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 27, 2014, pp. 3320–3328.
- [20] S. Kornblith, J. Shlens, and Q. V. Le, "Do better ImageNet models transfer better?" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 2661–2671.
- [21] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Front. Plant Sci.*, vol. 7, Art. no. 1419, 2016.
- [22] J. G. A. Barbedo, "Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification," *Comput. Electron. Agric.*, vol. 153, pp. 46–53, 2018.
- [23] M. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep learning for tomato diseases: Classification and symptoms visualization," *Appl. Artif. Intell.*, vol. 31, no. 4, pp. 299–315, 2017.
- [24] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 618–626.
- [25] D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv:1511.08060, 2015.