

AI-Assisted Image-Based Lichen Health Assessment for Environmental Sensing Using Naturally Acquired Field Images

Andrew Wang¹ and Wei Li^{2,*}

¹ Geneva Community High School, Geneva, IL 60510, USA

² Department of Electrical and Computing Engineering, University of Victoria, Victoria, BC V8W 2Y2, Canada

* Correspondence: weili@uvic.ca

How To Cite: Wang, A.; Li, W. AI-Assisted Image-Based Lichen Health Assessment for Environmental Sensing Using Naturally Acquired Field Images. *Journal of Artificial Intelligence for Automation* **2026**, *1*(3), 16. <https://doi.org/10.53941/jaia.2026.100016>

Received: 14 July 2026

Revised: 17 September 2026

Accepted: 20 September 2026

Published: 15 October 2026

Abstract: Lichens are widely recognized as biological indicators of environmental quality because their physiological condition reflects long-term atmospheric and ecological changes. However, conventional lichen assessment relies primarily on manual visual inspection by trained specialists, limiting the scalability of environmental monitoring. This paper presents an AI-assisted image analysis framework for automated lichen health assessment using naturally acquired field images as a step toward intelligent environmental sensing. A dataset consisting of 140 images of *Lobaria pulmonaria* was collected from the public iNaturalist database and manually annotated according to health and moisture conditions. An end-to-end framework integrating dataset construction, image augmentation, deep feature learning, health classification, and quantitative evaluation was developed and evaluated using stratified five-fold cross-validation. Binary health assessment achieved an average classification accuracy of 82.1%, while binary moisture and four-class lichen-status classification achieved accuracies of 75.7% and 64.3%, respectively. To further evaluate the robustness of the proposed framework, representative convolutional neural network, transformer-based, and classical machine learning classifiers were compared under an identical experimental protocol, with EfficientNet-B0 achieving the highest overall performance. The results demonstrate that AI can reliably assess lichen health from naturally acquired field images and further reveal that environmental moisture introduces substantial visual ambiguity for automated health assessment. This work establishes a practical and reproducible baseline for future AI-assisted environmental sensing, ecological monitoring, and image-based environmental bioindication.

Keywords: machine learning; image classification; lichen health assessment; environmental bioindication; computer vision

1. Introduction

Environmental monitoring plays an essential role in protecting ecosystem sustainability, improving public health, and supporting informed environmental management. Among various biological indicators, lichens have long been recognized as valuable bioindicators because their growth, morphology, and physiological condition are closely influenced by long-term atmospheric conditions and environmental stressors, including sulfur dioxide, nitrogen oxides, heavy metals, and particulate matter [1]. Owing to their high sensitivity to environmental change, many epiphytic lichen species provide an integrated biological response to air quality over extended periods, making them effective indicators of ecosystem health and environmental quality.

Traditional lichen assessment relies primarily on manual visual inspection performed by trained specialists. Field surveys typically require experts to identify lichen species, evaluate physiological condition, and estimate coverage before environmental indices can be derived [2]. Although these approaches are well established and scientifically

reliable, they remain labor-intensive, time-consuming, and difficult to scale for large geographic regions or repeated long-term monitoring. Developing automated image-based assessment techniques therefore represents an important step toward practical environmental sensing and ecological monitoring. Recent sensing technologies have also demonstrated the feasibility of in-field optical assessment of biological conditions using low-cost imaging systems [3].

Recent advances in artificial intelligence (AI), particularly deep learning and computer vision, have transformed automated image analysis in agriculture, ecology, environmental science, and biomedical applications [4–6]. AI-based image classification has achieved remarkable performance in plant recognition, biodiversity assessment, and ecological monitoring [2,7,8]. Existing lichen-related studies have primarily focused on species identification, instance segmentation, biodiversity assessment, or large-scale lichen coverage estimation using field photographs or remote sensing imagery [2,7,9]. These studies clearly demonstrate the potential of AI-assisted ecological monitoring while highlighting the growing importance of intelligent environmental sensing technologies.

Despite these advances, comparatively little attention has been devoted to automated assessment of lichen health from naturally acquired field images. Unlike species identification, physiological health assessment requires distinguishing subtle morphological and color changes associated with environmental stress while simultaneously accounting for natural variations introduced by moisture, illumination, viewing angle, background complexity, and image quality. These factors substantially increase the difficulty of image-based health assessment and remain important challenges for practical environmental bioindication.

To address these challenges, this study investigates the feasibility of an AI-assisted image analysis pipeline for automated lichen health assessment using naturally acquired images of *Lobaria pulmonaria*. Rather than proposing a new deep learning architecture, the objective is to evaluate whether existing AI techniques can provide reliable image-based environmental sensing under realistic field conditions. The proposed framework integrates dataset construction, image augmentation, deep feature learning, health classification, and quantitative performance evaluation into a unified image analysis pipeline.

The primary contributions of this work are threefold. First, an AI-assisted image analysis pipeline is developed for automated lichen health assessment using naturally acquired field images. Second, the influence of moisture on automated health assessment is systematically investigated through both four-class lichen-status classification and binary health and moisture classification tasks. Third, representative AI classification strategies, including convolutional neural networks, transformer-based models, and classical machine learning classifiers, are evaluated under an identical stratified five-fold cross-validation framework to investigate the robustness of the proposed environmental sensing approach. The experimental results establish a practical baseline for future AI-assisted environmental sensing, ecological monitoring, and image-based bioindication, as shown in Figure 1.




Figure 1. Overview of the proposed image-based lichen health assessment framework. Field images collected from the iNaturalist database are processed through dataset construction, feature extraction, health assessment, and performance evaluation to investigate the feasibility of automated lichen health assessment as an initial step toward environmental sensing.

2. Related Work and Research Gap

Lichens have long been recognized as valuable biological indicators because they respond sensitively to long-term environmental stress, including atmospheric pollutants, climate variation, and habitat quality [1,2,10]. Traditional lichen-based environmental assessment relies on field surveys performed by trained specialists, who identify lichen species and evaluate their abundance, coverage, or physiological condition to estimate environmental quality. Although these approaches have proven effective, they are labor-intensive, time-consuming, and difficult to scale for large geographic regions or repeated environmental monitoring.

Recent advances in computer vision and machine learning have enabled automated analysis of biological images for ecological applications. Existing lichen-related studies have primarily focused on species identification, biodiversity assessment, instance segmentation, or large-scale lichen coverage estimation. For example, Pärtel et al. [11] demonstrated high classification accuracy for automated plant identification using deep learning. Presta et al. [7] developed an image-based framework for automatic lichen species classification, while Nimi et al. [2] compared convolutional neural networks and classical image descriptors for lichen recognition. Kennedy et al. [9] further extended machine learning to satellite-based estimation of terricolous lichen cover for regional ecological studies. Collectively, these studies demonstrate the growing potential of AI-assisted image analysis for ecological monitoring and environmental sensing. Recent studies have demonstrated that advanced representation learning can support image classification across a range of environmental and sensing applications. Zhang et al. [12] proposed GCU-Net, which incorporates geospatial cognition into coral-reef geomorphology classification. Ye et al. [13] combined multimodal remote-sensing information with feature classification for pine wilt disease monitoring. These studies demonstrate the value of incorporating environmental context and complementary sensing information into ecological classification. In contrast, the present study focuses on close-range RGB field photographs collected under uncontrolled conditions and explicitly investigates moisture as a source of visual ambiguity in lichen health assessment, as shown in Table 1.

Table 1. Representative AI-based studies relevant to lichen and ecological image analysis.

Reference	Target	Image Source	Primary Objective	Representative AI Method	Main Contribution
Pärtel et al. [11]	Multiple plant species	Field photographs	Plant species identification	CNN + location information	Demonstrated high-accuracy automated plant identification under field conditions.
Naimi et al. [2]	20 lichen species	Time-lapse images	Species recognition	CNN + SIFT descriptors	Compared deep learning and classical image descriptors for automated lichen recognition.
Kennedy et al. [9]	Lichen cover	Satellite imagery	Regional ecological monitoring	Neural network	Estimated lichen cover over large geographic regions using remote sensing.
Presta et al. [7]	Lichens	Digital images	Automatic lichen classification	Deep learning	Developed an image-based framework for automated lichen classification.
This work	<i>Lobaria pulmonaria</i>	iNaturalist field images	Health assessment for environmental sensing	CNN, Transformer, SVM, RF	Investigates automated lichen health assessment and systematically compares representative AI classification strategies under an identical experimental protocol.

Despite these advances, comparatively little attention has been devoted to automated assessment of lichen physiological health from naturally acquired field images. Unlike species identification, health assessment requires distinguishing subtle morphological and color changes caused by environmental stress while simultaneously accounting for natural variations introduced by moisture, illumination, viewing angle, background complexity, and image quality. These factors substantially increase the difficulty of image-based health assessment and remain important challenges for practical environmental bioindication.

Furthermore, most previous studies employ a single machine learning model to address a specific ecological task. Few studies have investigated whether different AI classification strategies produce consistent health assessment results under identical experimental conditions. Such comparative evaluation is important for understanding the robustness, generalization, and practical applicability of AI-assisted environmental sensing systems.

Motivated by these observations, this study investigates image-based lichen health assessment as an important step toward automated environmental sensing. A curated dataset of *Lobaria pulmonaria* field images was constructed from the public iNaturalist database and manually labeled according to both health and moisture conditions. Representative deep-learning, transformer-based, and classical machine learning classifiers are evaluated using

an identical stratified five-fold cross-validation protocol to enable a fair comparison of different AI strategies. The resulting framework establishes a practical baseline for future AI-assisted environmental sensing, ecological monitoring, and image-based bioindication, as shown in Table 2.

Table 2. Comparison of representative studies. The (✓) indicate 'Yes' and (–) indicate 'No'.

Study	Species	Health	Moisture	Env. Sensing	AI Comparison
Pärtel et al. [11]	✓	–	–	–	–
Naimi et al. [2]	✓	–	–	–	✓
Kennedy et al. [9]	–	–	–	✓	–
Presta et al. [7]	✓	–	–	✓	–
This work	–	✓	✓	✓	✓

3. Materials and Methods

3.1. Overall Framework

The proposed AI-assisted environmental sensing framework is illustrated in Figure 1. The framework consists of five sequential stages:

- (1) Image acquisition;
- (2) Dataset construction;
- (3) AI-based feature learning;
- (4) Lichen health classification;
- (5) Performance evaluation.

The objective is to investigate whether naturally acquired lichen images can provide reliable information for automated health assessment and environmental sensing. Rather than developing a new deep learning architecture, this study establishes a complete image analysis pipeline and evaluates multiple representative AI classification strategies under an identical experimental protocol.

Figure 2 illustrates the experimental framework used to compare four representative AI classification strategies. Two end-to-end deep learning models, EfficientNet-B0 and MobileViT-S, directly perform image classification using a fully connected layer followed by a Softmax classifier. In addition, the 1280-dimensional feature vectors extracted by EfficientNet-B0 are used as inputs to two classical machine learning classifiers, namely Support Vector Machine (SVM) and Random Forest [14–17]. This design enables systematic evaluation of different decision mechanisms while maintaining a common feature representation, thereby providing a fair comparison of representative deep learning, transformer-based, and classical machine learning approaches for automated lichen health assessment.




Figure 2. Experimental framework for comparative evaluation of four representative AI classification strategies. EfficientNet-B0 is evaluated both as an end-to-end convolutional neural network with a Softmax classifier and as a

deep feature extractor whose learned feature vectors are classified using Support Vector Machine (SVM) and Random Forest (RF). MobileViT-S is independently evaluated as a lightweight CNN-transformer architecture with an end-to-end Softmax classifier. All models are trained and evaluated under the same experimental protocol for binary health assessment, binary moisture assessment, and four-class lichen-status classification.

3.2. Image Dataset

This study investigates automated lichen health assessment using *Lobaria pulmonaria*, an epiphytic lichen species widely recognized as an environmental bioindicator.

Images were collected from the public iNaturalist database. Only observations labeled as *Research Grade* were included to improve species identification reliability. Unlike laboratory datasets, these field photographs exhibit substantial variations in illumination, viewing angle, background complexity, camera quality, and environmental conditions, providing a realistic evaluation environment for AI-assisted environmental sensing.

All images were manually reviewed before inclusion. Images containing severe blur, significant occlusion, multiple dominant specimens, or insufficient visibility were discarded.

Each image was manually labeled according to both physiological condition and moisture status, resulting in four lichen-status classes:

- Healthy–Wet (HW)
- Healthy–Dry (HD)
- Unhealthy–Wet (UW)
- Unhealthy–Dry (UD)

The final dataset contains 140 annotated field images. The class distribution was relatively balanced, comprising 36 Healthy–Dry images, 33 Healthy–Wet images, 35 Unhealthy–Dry images, and 36 Unhealthy–Wet images. Consequently, the binary health dataset contained 69 healthy and 71 unhealthy images, while the binary moisture dataset contained 71 dry and 69 wet images.

3.2.1. Annotation Criteria

Health and moisture labels were assigned independently based on visual examination of each image. A thallus was labeled as *unhealthy* when at least two of the following indicators were clearly visible: (1) substantial discoloration or bleaching; (2) localized or widespread necrotic brown or black patches; (3) fragmentation, damaged lobe margins, or lobe loss; (4) exposure of the pale medulla resulting from erosion of the pigmented surface layer; and (5) persistent loss of turgidity, including a shriveled, curled, or brittle appearance that could not be explained solely by dry conditions. Specimens exhibiting characteristic green–gray coloration, intact lobe margins, and a relatively uniform and continuous thallus surface were labeled as *healthy*. Borderline samples showing only mild or isolated discoloration were conservatively assigned to the healthy category to avoid interpreting normal pigmentation variation as tissue damage.

Moisture status was evaluated separately from health condition. A thallus was labeled as *wet* when it exhibited darker and more saturated coloration together with smooth, hydrated-looking lobes. A thallus was labeled as *dry* when it exhibited lighter coloration and a comparatively brittle, wrinkled, or curled appearance. These health and moisture criteria produced four classes: Healthy–Wet (HW), Healthy–Dry (HD), Unhealthy–Wet (UW), and Unhealthy–Dry (UD). All images were labeled by a single annotator using the visual criteria described above. No laboratory, chemical, or longitudinal measurements were available to confirm the physiological condition of the photographed specimens. Accordingly, the labels represent visually assessed health conditions rather than physiologically confirmed diagnoses. Inter-annotator agreement could not be evaluated, which represents a limitation of the current dataset.

3.2.2. Deep Feature Learning

EfficientNet-B0 was adopted as the primary feature learning network because it provides an effective balance between classification performance, computational efficiency, and model complexity for relatively small image datasets.

Given an input image $\mathbf{I}$,

$$\mathbf{f} = F_{\theta}(\mathbf{I}), \quad (1)$$

where $F_{\theta}(\cdot)$ denotes the pretrained EfficientNet-B0 network and $\mathbf{f}$ represents the learned 1280-dimensional feature vector.

Global Average Pooling converts the final feature maps into compact image descriptors,

$$f_c = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W x_{ijc}. \quad (2)$$

Posterior probabilities are obtained using

$$P(y = k) = \frac{\exp(z_k)}{\sum_j \exp(z_j)}, \quad (3)$$

and optimized using cross-entropy loss

$$L = - \sum_k y_k \log P(y = k). \quad (4)$$

3.2.3. Classification Strategies

To investigate the robustness of the learned image representation, four representative AI classification strategies were evaluated: EfficientNet-B0 with Softmax classification, MobileViT-S, EfficientNet-B0 features with Support Vector Machine (SVM), and EfficientNet-B0 features with Random Forest (RF). These methods represent convolutional neural networks, lightweight transformer-based models, and classical machine learning classifiers.

EfficientNet-B0 with Softmax

EfficientNet-B0 was used as the primary convolutional feature learning model. Given an input image $\mathbf{I}$, the pretrained backbone extracts a compact feature representation,

$$\mathbf{f} = F_\theta(\mathbf{I}), \quad (5)$$

where $F_\theta(\cdot)$ denotes the EfficientNet-B0 feature extractor. The learned feature vector is expressed as

$$\mathbf{f} = [f_1, f_2, \dots, f_{1280}]^T, \quad (6)$$

where each element represents a learned latent image feature. These features collectively encode visual information related to color, texture, morphology, and structural appearance.

The final classification layer maps the feature vector to class scores,

$$\mathbf{z} = \mathbf{W}\mathbf{f} + \mathbf{b}, \quad (7)$$

where $\mathbf{W}$ and $\mathbf{b}$ are the classifier weights and bias terms. The posterior probability of class k is obtained using the Softmax function,

$$P(y = k) = \frac{\exp(z_k)}{\sum_{j=1}^K \exp(z_j)}. \quad (8)$$

The predicted label is then determined by

$$\hat{y} = \arg \max_k P(y = k). \quad (9)$$

During training, the network parameters were optimized using cross-entropy loss,

$$L = - \sum_{k=1}^K y_k \log P(y = k), \quad (10)$$

where y_k is the ground-truth label indicator.

Support Vector Machine

For the SVM classifier, the EfficientNet-B0 backbone was used as a fixed feature extractor, and the resulting 1280-dimensional feature vectors were used as the input to the SVM [17]. The SVM seeks an optimal separating hyperplane by solving

$$\min_{\mathbf{w}, b, \xi} \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^N \xi_i, \quad (11)$$

subject to

$$y_i(\mathbf{w}^T \mathbf{x}_i + b) \geq 1 - \xi_i, \quad (12)$$

where $\mathbf{x}_i$ is the feature vector of the i th image, $y_i \in \{-1, +1\}$ is the class label, ξ_i is the slack variable, and C controls the trade-off between margin maximization and classification error.

An RBF kernel was used to allow nonlinear decision boundaries in the learned feature space,

$$K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2), \quad (13)$$

where γ controls the kernel width. The SVM prediction is given by

$$\hat{y} = \text{sign} \left(\sum_{i=1}^{N_s} \alpha_i y_i K(\mathbf{x}_i, \mathbf{x}) + b \right), \quad (14)$$

where N_s is the number of support vectors and α_i are the learned support-vector coefficients. In this study, SVM classification evaluates whether the deep features extracted by EfficientNet-B0 are separable using a classical maximum-margin decision boundary.

Random Forest

For the Random Forest classifier, the same EfficientNet-B0 feature vectors were used as input [16]. A Random Forest consists of multiple decision trees, each of which produces an individual class prediction,

$$T_m(\mathbf{x}) = h_m(\mathbf{x}), \quad (15)$$

where $T_m(\cdot)$ is the m th decision tree and $h_m(\cdot)$ denotes its learned decision function. The final prediction is obtained by majority voting across all trees,

$$\hat{y} = \text{mode} \{T_1(\mathbf{x}), T_2(\mathbf{x}), \dots, T_M(\mathbf{x})\}. \quad (16)$$

The class probability can be estimated as the proportion of trees voting for class c ,

$$P(y = c) = \frac{1}{M} \sum_{m=1}^M \mathbb{I}(T_m(\mathbf{x}) = c), \quad (17)$$

where $\mathbb{I}(\cdot)$ is the indicator function and M is the number of trees. This classifier evaluates whether the learned feature space can be separated using an ensemble of threshold-based decision rules.

MobileViT-S

MobileViT-S was evaluated as a representative lightweight transformer-based image classification model [15,18]. Unlike EfficientNet-B0, which relies primarily on convolutional feature extraction, MobileViT-S combines local convolutional representations with transformer-based global context modeling. Conceptually, its feature representation can be written as

$$\mathbf{F}_{\text{MobileViT}} = \mathcal{F}(\mathbf{F}_{\text{CNN}}, \mathbf{F}_{\text{Transformer}}), \quad (18)$$

where $\mathbf{F}_{\text{CNN}}$ represents local convolutional features, $\mathbf{F}_{\text{Transformer}}$ represents global contextual features, and $\mathcal{F}(\cdot)$ denotes their integration within the MobileViT architecture. MobileViT-S was included to compare the proposed CNN-based approach with a recent lightweight transformer-based classifier under the same experimental protocol.

3.2.4. Image Augmentation

To improve model generalization under realistic field conditions, several image augmentation techniques were applied during training, including random resized cropping, horizontal and vertical flipping, color jittering, grayscale conversion, Gaussian blur, and random erasing.

MixUp augmentation was introduced during the second training stage,

$$\tilde{x} = \lambda x_i + (1 - \lambda) x_j, \quad (19)$$

where $\lambda \sim \text{Beta}(\alpha, \alpha)$.

The corresponding labels become

$$\tilde{y} = \lambda y_i + (1 - \lambda) y_j. \quad (20)$$

These augmentation strategies improve robustness against illumination variation, viewpoint changes, partial occlusion, and the limited dataset size.

3.3. Experimental Setup

All images were converted to RGB format and normalized using the ImageNet mean and standard deviation. During validation, each image was resized to 256 pixels and center-cropped to 224×224 pixels. Training augmentation included random resized cropping, horizontal and vertical flipping, color jittering, grayscale conversion, Gaussian blurring, and random erasing. Model performance was evaluated using stratified five-fold cross-validation on the complete dataset of 140 images, with approximately 112 images used for training and 28 images used for validation in each fold.

ImageNet-pretrained EfficientNet-B0 and MobileViT-S models were trained using a two-stage transfer-learning procedure. In the first stage, the pretrained backbone was frozen and only the classification head was trained for five epochs using a learning rate of 1×10^{-3} . In the second stage, the entire network was fine-tuned for 20 epochs using a learning rate of 1×10^{-4} . AdamW was used for optimization with a weight decay of 0.05 and a cosine-annealing learning-rate scheduler with $T_{\max} = 20$. The batch size was set to 8, and cross-entropy was used as the loss function. For EfficientNet-B0, a dropout rate of 0.4 was applied in the classification head. A recent review of the evolution and roles of deep-learning optimizers is provided by Wen and Zhou [19], offering broader methodological context for the selection of adaptive optimization algorithms in image-classification tasks. MixUp augmentation with $\alpha = 0.2$ was applied to 50% of the fine-tuning batches. During evaluation, test-time augmentation was implemented by averaging the prediction probabilities obtained from five transformed views.

For the classical machine-learning models, the pretrained EfficientNet-B0 backbone was frozen and used to extract a 1,280-dimensional feature vector from each image. For Support Vector Machine classification, the extracted features were standardized within each training fold, and an RBF kernel was employed with $C = 10$ and balanced class weights. Platt scaling was used to obtain calibrated probability estimates. The Random Forest classifier consisted of 300 decision trees, used balanced class weights, and imposed no maximum tree depth. Feature scaling was not applied to the Random Forest classifier.

Model evaluation was performed using stratified five-fold cross-validation to preserve class distribution across all folds. Test-Time Augmentation (TTA) was employed by averaging prediction probabilities obtained from five transformed views,

$$P = \frac{1}{M} \sum_{m=1}^M P_m. \quad (21)$$

Classification performance was evaluated using accuracy, precision, recall, specificity, F1-score, and confusion matrices.

For binary health assessment, the four lichen-status classes were further grouped into healthy and unhealthy categories,

$$g(y) = \begin{cases} \text{Healthy}, & y \in \{\text{HW}, \text{HD}\}, \\ \text{Unhealthy}, & y \in \{\text{UW}, \text{UD}\}. \end{cases} \quad (22)$$

Both direct binary classification and aggregation from the four-class classifier were investigated to evaluate the influence of moisture on automated lichen health assessment.

3.4. Performance Metrics

Classification performance was evaluated using accuracy, precision, recall (sensitivity), specificity, F1-score, and confusion matrices. For binary health assessment, unhealthy lichen was treated as the positive class and healthy lichen as the negative class. The evaluation metrics were calculated as

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (23)$$

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (24)$$

$$\text{Recall} = \frac{TP}{TP + FN}, \quad (25)$$

$$\text{Specificity} = \frac{TN}{TN + FP}, \quad (26)$$

$$F_1 = \frac{2TP}{2TP + FP + FN}, \quad (27)$$

where TP , TN , FP , and FN denote the numbers of true positives, true negatives, false positives, and false negatives, respectively. Metrics were calculated separately for each cross-validation fold and are reported as mean $\pm$ standard deviation. In addition, predictions from all five held-out folds were pooled to construct an aggregate confusion matrix.

4. Results and Discussion

The proposed AI-assisted environmental sensing framework was evaluated using stratified five-fold cross-validation on the complete dataset consisting of 140 naturally acquired field images. Three progressively more challenging classification tasks were investigated, including binary health assessment, binary moisture assessment, and four-class lichen-status classification. The reported results correspond to the best validation model obtained within each cross-validation fold.

Figure 3 presents representative validation examples from the best-performing cross-validation fold. Correctly classified samples are observed across all four lichen-status categories, demonstrating that the proposed AI-assisted framework successfully captures meaningful physiological characteristics from naturally acquired field images. Consistent with the aggregated confusion matrix, most misclassifications occur between healthy-dry and unhealthy-dry specimens, whereas healthy-wet samples are generally recognized with high confidence. These examples illustrate the considerable visual similarity among certain dry specimens and further support the observation that environmental moisture substantially influences automated lichen health assessment.




Figure 3. Representative validation images from the best-performing cross-validation fold. Green borders indicate correctly classified samples, while red borders denote misclassified samples. The figure includes examples from all four lichen-status classes: healthy-dry (HD), healthy-wet (HW), unhealthy-dry (UD), and unhealthy-wet (UW). For each image, the ground-truth label, predicted label, and prediction confidence are shown. Correctly classified healthy-wet and unhealthy-dry specimens demonstrate the effectiveness of the proposed framework, whereas most misclassifications occur between healthy-dry and unhealthy-dry samples, highlighting the visual ambiguity introduced by dry environmental conditions.

4.1. Binary Health Assessment

The primary objective of this study was to investigate whether the proposed framework could distinguish healthy and unhealthy *Lobaria pulmonaria* from naturally acquired field images. Binary health assessment achieved

a mean five-fold classification accuracy of $82.1\% \pm 7.5\%$. The accuracies for the five folds were 75.0%, 96.4%, 78.6%, 78.6%, and 78.6%, respectively, as shown in Table 3. The observed fold-to-fold variability reflects the limited dataset size, because each held-out fold contained only approximately 28 images.

Table 3. Comparison of representative AI classification strategies. The best-performing metrics are highlighted in bold.

Classifier	Feature Learning	Decision Strategy	Binary Health (%)	Binary Moisture (%)	4-Class (%)	Characteristics
EfficientNet-B0	CNN	Softmax	82.1	75.7	64.3	Highest overall accuracy; end-to-end deep feature learning.
MobileViT-S	CNN +Transformer	Softmax	79.3	75.0	–	Competitive performance with lightweight hybrid architecture.
EfficientNet + RF	CNN	Random Forest	72.1	72.9	–	Ensemble decision trees operating on learned deep features.
EfficientNet + SVM	CNN	RBF SVM	69.3	65.0	–	Maximum-margin classifier using EfficientNet feature vectors.

Across the five cross-validation folds, the best validation accuracies were 75.0%, 96.4%, 82.1%, 78.6%, and 82.1%, respectively. Although moderate variation exists between folds, this behavior is expected for a relatively small dataset containing only 140 images.

The two-stage transfer learning strategy consistently improved performance after fine-tuning the pretrained EfficientNet-B0 backbone. The results indicate that transfer learning provides effective feature representations even when only a limited number of ecological images are available. This observation demonstrates that transferring visual representations learned from large-scale natural image datasets provides an effective starting point for ecological image analysis, even when only a limited number of labeled lichen images are available.

As shown in Table 4, the model achieved a mean sensitivity of $85.9 \pm 10.1\%$ for unhealthy specimens and a mean specificity of $78.2 \pm 13.6\%$ for healthy specimens. The aggregate confusion matrix, presented in Table 5, contained 61 true-positive, 54 true-negative, 15 false-positive, and 10 false-negative predictions. These counts produced an aggregate precision of 80.3%, recall of 85.9%, specificity of 78.3%, and F1-score of 83.0%, as shown in Table 6. The stronger recognition of unhealthy specimens suggests that visible tissue degradation and discoloration may provide more distinctive visual cues than those present in healthy specimens.

Table 4. Five-fold performance of EfficientNet-B0 for binary health assessment. Unhealthy lichen was treated as the positive class.

Fold	Accuracy (%)	Precision (%)	Recall (%)	Specificity (%)	F1 (%)
1	75.0	76.9	71.4	78.6	74.1
2	96.4	100.0	92.9	100.0	96.3
3	78.6	78.6	78.6	78.6	78.6
4	78.6	70.0	100.0	57.1	82.4
5	82.1	81.2	86.7	76.9	83.9
Mean	82.1	81.3	85.9	78.2	83.0
SD	7.5	10.0	10.1	13.6	7.4

Table 5. Aggregate confusion matrix for binary health assessment across the five held-out folds.

Actual Class	Predicted Class	
	Healthy	Unhealthy
Healthy	54	15
Unhealthy	10	61

Table 6. Performance calculated from the aggregate binary health confusion matrix.

Metric	Value (%)
Accuracy	82.1
Precision	80.3
Recall	85.9
Specificity	78.3
F1-score	83.0

4.2. Influence of Moisture on Health Assessment

Although binary health assessment achieved encouraging performance, environmental moisture substantially influences lichen appearance and therefore represents an important source of classification uncertainty.

Binary moisture classification achieved an average accuracy of $75.7 \pm 9.1\%$, lower than binary health classification. Per-class accuracies of $78.9 \pm 4.6\%$ for dry specimens and $73.8 \pm 18.5\%$ for wet specimens further indicate larger performance variation under different environmental conditions.

These observations suggest that moisture introduces considerable visual variability, including changes in color, texture, and surface reflectance, partially masking the physiological characteristics associated with lichen health. These findings suggest that physiological health and moisture are not completely independent visual attributes. Instead, changes in moisture alter the appearance of otherwise healthy and unhealthy specimens, thereby increasing intra-class variability and reducing class separability. Consequently, moisture represents one of the major challenges for practical image-based environmental sensing.

4.3. Toward Comprehensive Lichen Status Assessment

To further evaluate the proposed framework under a more challenging and realistic scenario, four-class lichen-status classification was investigated by simultaneously considering both physiological condition and moisture state.

Compared with binary health assessment, four-class classification achieved a lower average accuracy of $64.3 \pm 4.5\%$. This reduction is expected because the classifier must simultaneously distinguish multiple sources of visual variation. The four-class classification results provide important insight into the practical limits of current image-based environmental assessment. While binary health assessment is already suitable for preliminary environmental screening, reliable simultaneous recognition of health and moisture requires substantially more discriminative visual representations and larger annotated datasets.

Inspection of the aggregated confusion matrix indicates that most classification errors occurred between healthy-dry and unhealthy-dry specimens, whereas healthy-wet specimens were recognized more consistently. These observations further support the hypothesis that moisture significantly influences visual appearance and increases the difficulty of automated health assessment.

Rather than indicating a limitation of the proposed AI framework, these results demonstrate that simultaneous recognition of physiological condition and environmental moisture remains a substantially more difficult environmental sensing problem. Consequently, four-class classification naturally motivates future research toward more comprehensive AI-assisted ecological monitoring systems.

4.4. Comparison of Representative AI Classification Strategies

To evaluate the robustness of the proposed environmental sensing framework, representative deep learning, transformer-based, and classical machine learning classifiers were evaluated using an identical experimental protocol.

EfficientNet-B0 achieved the highest overall classification accuracy for both binary health and moisture assessment. Nevertheless, MobileViT-S demonstrated competitive performance, indicating that reliable lichen health assessment is not restricted to a single neural network architecture. The relatively small performance gap between EfficientNet-B0 and MobileViT-S further indicates that recent lightweight vision transformer architectures are also promising candidates for future ecological image analysis.

Although Random Forest and Support Vector Machine produced lower overall accuracies, both methods successfully classified lichen health with reasonable performance. These findings suggest that the learned image representations contain meaningful physiological information that can be exploited by both deep learning and classical machine learning algorithms. The relatively low computational complexity of these classifiers may also make them attractive for preliminary image screening or resource-constrained environmental monitoring applications.

Overall, the experimental results demonstrate that the proposed AI-assisted image analysis framework provides a practical and promising solution for automated lichen health assessment using naturally acquired field images. More importantly, the progression from binary health assessment to moisture recognition and finally four-class lichen-status classification illustrates increasing levels of problem complexity. While reliable health assessment has been demonstrated, simultaneous interpretation of multiple physiological and environmental attributes remains an open research challenge. The present study establishes a baseline for future research on hierarchical environmental sensing, multi-task learning, explainable AI, and large-scale ecological monitoring using naturally acquired field images [20, 21]. Future work could investigate more flexible learning settings as the lichen dataset expands. For example, exemplar-free class-incremental learning [22] could allow newly annotated health conditions or lichen species to be incorporated without retaining all previously collected images. Cross-modal and open-set

representation-learning methods [23] may also support the integration of RGB photographs with environmental variables while improving the recognition of samples outside the known training categories.

5. Conclusions

This paper presents an AI-assisted image analysis framework for automated lichen health assessment using naturally acquired field images of *Lobaria pulmonaria*. A complete environmental sensing pipeline, including dataset construction, image augmentation, AI-based feature learning, health classification, and quantitative performance evaluation, has been developed to investigate the feasibility of image-based lichen health assessment under realistic field conditions.

Experimental results demonstrate that the proposed framework can reliably distinguish healthy and unhealthy lichen specimens using naturally acquired images. Binary health assessment achieves the highest classification performance, while additional experiments show that environmental moisture introduces substantial visual variability that increases the difficulty of automated health assessment. Furthermore, comparison with representative deep learning, transformer-based, and classical machine learning classifiers demonstrates that the proposed framework maintains consistent performance across different AI paradigms, with EfficientNet-B0 achieving the highest overall classification accuracy.

The progression from binary health assessment to moisture recognition and finally four-class lichen-status classification illustrates increasing levels of problem complexity for AI-assisted environmental sensing. These results suggest that visually assessed health status can be distinguished with encouraging accuracy from field images, whereas simultaneous interpretation of both health condition and environmental moisture remains a more challenging research problem.

Although this study considers a single lichen species and a relatively small image dataset, it establishes a practical and reproducible baseline for AI-assisted environmental sensing using biological indicators. Future research will investigate larger multi-species datasets, hierarchical and multi-task classification strategies, integration of environmental metadata, and temporal image analysis to further advance automated ecological monitoring and image-assisted environmental bioindication. By transforming biological observations into scalable digital information, AI-assisted environmental sensing has the potential to support intelligent automation for long-term ecosystem monitoring and evidence-based environmental management [24].

Author Contributions

W.L.: Conceptualization, methodology, supervision, and writing—review and editing. A.W.: Methodology, software, investigation, and writing—original draft preparation. Both authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The original images used in this study were obtained from the publicly accessible iNaturalist database (<https://www.inaturalist.org/>). The curated annotations and supporting data generated during this study are available from the corresponding author upon reasonable request.

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

During the preparation of this manuscript, the authors used ChatGPT (OpenAI) for language editing, proofreading, and improving readability. All AI-assisted changes were carefully reviewed and verified by the authors, who take full responsibility for the final content.

References

1. Glavich, D.A.; Geiser, L.H. Potential Approaches to Developing Lichen-Based Critical Loads and Levels for Nitrogen, Sulfur and Metal-Containing Atmospheric Pollutants in North America. *Bryologist* **2008**, *111*, 638–649. <https://doi.org/10.1639/0007-2745-111.4.638>.
2. Nimis, P.L.; Scheidegger, C.; Wolseley, P.A. *Monitoring with Lichens—Monitoring Lichens*; Springer: Dordrecht, Netherlands, 2002. <https://doi.org/10.1007/978-94-010-0423-7>.
3. Lauretti, C.; Tamantini, C.; Zompanti, A.; et al. A Low-Cost Multispectral Device for in-Field Fruit Ripening Assessment. *IEEE Sens. J.* **2025**, *25*, 27834–27844. <https://doi.org/10.1109/jsen.2023.3335259>.
4. Perry, G.L.W.; Seidl, R.; Bellvé, A.M.; et al. An Outlook for Deep Learning in Ecosystem Science. *Ecosystems* **2022**, *25*, 1700–1718. <https://doi.org/10.1007/s10021-022-00789-y>.
5. Trotter, C.; Griffiths, H.J.; Whittle, R.J. Surveying the Deep: A Review of Computer Vision in the Benthos. *Ecol. Inform.* **2025**, *86*, 102989. <https://doi.org/10.1016/j.ecoinf.2024.102989>.
6. Butera, L.; Ferrante, A.; Jermini, M.; et al. Precise Agriculture: Effective Deep Learning Strategies to Detect Pest Insects. *IEEE/CAA J. Autom. Sin.* **2022**, *9*, 246–258. <https://doi.org/10.1109/jas.2021.1004317>.
7. Presta, A.; Pellegrino, F.A.; Martellos, S. Learning-Based Automatic Classification of Lichens from Images. *Biosyst. Eng.* **2022**, *213*, 119–132. <https://doi.org/10.1016/j.biosystemseng.2021.11.023>.
8. Tuia, D.; Kellenberger, B.; Beery, S.; et al. Perspectives in Machine Learning for Wildlife Conservation. *Nat. Commun.* **2022**, *13*, 792. <https://doi.org/10.1038/s41467-022-27980-y>.
9. Kennedy, B.; Pouliot, D.; Manseau, M.; et al. Assessment of Landsat-Based Terricolous Macrolichen Cover Retrieval and Change Analysis over Caribou Ranges in Northern Canada and Alaska. *Remote Sens. Environ.* **2020**, *240*, 111694. <https://doi.org/10.1016/j.rse.2020.111694>.
10. Nash, T.H., III. *Lichen Biology*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2008.
11. Pärtel, J.; Pärtel, M.; Wäldchen, J. Plant Image Identification Application Demonstrates High Accuracy in Northern Europe. *AoB PLANTS* **2021**, *13*, plab050. <https://doi.org/10.1093/aobpla/plab050>.
12. Zhang, F.; Hu, Y.; Ma, Y.; et al. GCU-Net: Remote Sensing Classification Method for Coral Reef Geomorphology Integrating Geospatial Cognition. *IEEE Geosci. Remote Sens. Lett.* **2024**, *21*, 5003805. <https://doi.org/10.1109/lgrs.2024.3406681>.
13. Ye, X.; Yu, H.; Yan, Y.; et al. Pine Wilt Disease Monitoring Using Multimodal Remote Sensing Data and Feature Classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2025**, *18*, 8536–8546. <https://doi.org/10.1109/jstars.2025.3549977>.
14. Tan, M.; Le, Q.V. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In Proceedings of the International Conference on Machine Learning (ICML), Long Beach, CA, USA, 9–15 June 2019.
15. Mehta, S.; Rastegari, M. MobileViT: Light-Weight, General-Purpose, and Mobile-Friendly Vision Transformer. In Proceedings of the International Conference on Learning Representations (ICLR), Virtual, 25–29 April 2022.
16. Breiman, L. Random Forests. *Mach. Learn.* **2001**, *45*, 5–32. <https://doi.org/10.1023/a:1010933404324>.
17. Cortes, C.; Vapnik, V. Support-Vector Networks. *Mach. Learn.* **1995**, *20*, 273–297. <https://doi.org/10.1023/a:1022627411411>.
18. Deng, J.; Dong, W.; Socher, R.; et al. ImageNet: A Large-Scale Hierarchical Image Database. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Miami Beach, FL, USA, 20–25 June 2009; pp. 248–255. <https://doi.org/10.1109/cvpr.2009.5206848>.
19. Wen, X.; Zhou, M. Evolution and Role of Optimizers in Training Deep Learning Models. *IEEE/CAA J. Autom. Sin.* **2024**, *11*, 2039–2042. <https://doi.org/10.1109/jas.2024.124806>.
20. Selvaraju, R.R.; Cogswell, M.; Das, A.; et al. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017; pp. 618–626. <https://doi.org/10.1109/iccv.2017.74>.
21. Lundberg, S.M.; Lee, S.I. A Unified Approach to Interpreting Model Predictions. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017.
22. Liu, M.; Ma, J.; Kang, Q.; et al. Exemplar-Free Class Incremental Learning for Traffic Sign Classification. *IEEE Trans. Veh. Technol.* **2026**, *75*, 3556–3568. <https://doi.org/10.1109/tvt.2025.3606740>.
23. Xiong, P.; Liao, J.; Zhou, M.; et al. Deeply Supervised Subspace Learning for Cross-Modal Material Perception of Known and Unknown Objects. *IEEE Trans. Ind. Inform.* **2023**, *19*, 2259–2268. <https://doi.org/10.1109/tii.2022.3195171>.
24. Zhou, M. Artificial Intelligence and Automation for Better Life of Human Beings. *J. Artif. Intell. Autom.* **2026**, *1*, 7. <https://doi.org/10.53941/jaia.2026.100007>.