Phyton-
International Journal of
Experimental Botany

Tech Science Press

http://doi.org/10.32604/phyton.2026.084983

ARTICLE

Expert-Based Phenotyping Using Predictive Modeling of Visual Assessments: Automated Severity Quantification of *Phytophthora infestans* in Potato Foliage

Nadezhda Onezhko*, Mariia Samarina, Aleksey Ermolaev, Daniil Ulyanov, Sergey Avdeev, Gennady Karlov and Mikhail Divashuk

Laboratory of Applied Genomics and Crop Breeding, All-Russia Research Institute of Agricultural Biotechnology, Moscow, Russia
*Corresponding Author: Nadezhda Onezhko. Email: one3.01@mail.ru

Received: 03 May 2026; Accepted: 30 July 2026; Published: 19 October 2026

ABSTRACT: A significant issue in phytopathological examinations is the subjectivity of visual assessments. This inter-observer variability leads to errors and reduces reproducibility, compromising the accuracy of research outcomes. Beyond human bias, the existence of numerous assessment scales, often inaccurately converted between one another, further complicates the problem. Therefore, automation of visual phytopathological evaluation is required. This study addresses the issue using the case of late blight infection in potato foliage. We propose a machine learning and computer vision pipeline that compares images of potato leaves against a standardized severity scale. Specifically, a convolutional neural network segments leaf images into regions, a support vector machine classifies infection levels, and an ensemble of decision trees approximates expert assessments across severity grades. The pipeline operates on visible symptoms in images acquired 3–4 days post-inoculation; pre-symptomatic or latent infection is outside its scope. The investigation was conducted using a dataset of potato leaf images affected by late blight, collected by the authors. The proposed methods achieved a pixel-wise classification accuracy of 92% (AUC = 0.974) and produced severity assessments that strongly correlated with reference expert evaluations (Pearson's $r = 0.821$, $p < 0.05$; Spearman's $\rho = 0.845$).

KEYWORDS: Computer vision; convolutional neural network (CNN); *Phytophthora infestans*; automated phenotyping; expert assessments

1 Introduction

Late blight, caused by the oomycete *Phytophthora infestans*, is one of the most destructive diseases of potato (*Solanum tuberosum* L.). Annual yield losses reach millions of tons worldwide, imposing severe economic burdens on potato-producing countries [1]. In response, substantial research efforts focus on developing resistant cultivars, chemical control strategies, and understanding the pathogen itself. Despite these investments, the disease remains a major threat.

A critical bottleneck in this research is the assessment of disease severity in plant material. Each new experiment requires repeated evaluation of infection levels, often performed visually by human experts. Visual assessment involves simultaneous scoring of multiple phenotypic traits [2]. This process is time-consuming, labor-intensive, and prone to inaccuracies. More importantly, it introduces subjectivity, which compromises experimental reproducibility—a cornerstone of reliable science. The problem is further exacerbated by the lack of a standardized scoring system across different research groups.

Copyright © 2026 The Authors. Published by Tech Science Press. This work is licensed under a Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

To address these challenges, automated image analysis tools offer a promising alternative. Such tools can process large numbers of plant images rapidly and provide quantitative, reproducible measurements. In this paper, we present an automated pipeline for quantifying late blight lesions on potato leaves and for predicting expert severity scores. Our approach first segments leaf images into individual lobes, then classifies pixels based on RGB channel brightness ratios, and finally calculates the percentage of infected area per lobe. This lobe-level processing is essential because standard resistance assessment protocols involve inoculating each lobe separately [3]. By aggregating lobe-level data, we predict the most likely expert assessment for the whole leaf.

The contribution of this work is threefold: (i) adaptation of an image-analysis and machine-learning pipeline to potato late blight; (ii) integration into an end-to-end workflow that approximates an expert ordinal score (1–9 scale); and (iii) release of a purpose-built annotated dataset. We do not claim novelty of the underlying machine-learning methods themselves. Our work is positioned as approximating an ordinal 1–9 expert (resistance) scale specifically for potato late blight, distinct from expert-validated continuous-severity pipelines on other crops.

We evaluate our approach by comparing model predictions against ground-truth expert assessments using correlation analysis (Pearson and Spearman) on a proprietary dataset of infected potato leaf images.

Existing Tools

One existing tool for annotating infected leaf areas is the Automatic Leaf Infection Identifier [4]. The software first detects whether an image contains any visible infection; for infected leaves, it draws a line along the main axis of each lesion. This is accomplished via a three-step CNN-based pipeline: (i) training CNNs to detect lesions in small image patches, (ii) generating infection probability maps for each image region, and (iii) classifying the original images from these maps.

However, when evaluated on our dataset, the tool failed to produce usable results. As shown in Fig. 1, it systematically misclassified entire leaf areas as infected while labeling healthy background tissue as healthy—effectively confusing leaf tissue with disease symptoms. This poor performance underscores the need for a specialized approach tailored to late blight in potato foliage.



Figure 1: Examples of Automatic-leaf-infection-identifier runs on the studied dataset, showing misclassification of leaf and background areas.

2 Materials and Methods

2.1 Plant Material and Infection Conditions

The plant material used in this study comprised mid-tier leaves detached from potato plants of the susceptible cultivar Vector (Solanum tuberosum L.), obtained from the in-house collection of the All-Russia Research Institute of Agricultural Biotechnology (Moscow, Russia), grown under controlled phytotron conditions, and inoculated with a zoospore suspension of the pathogen Phytophthora infestans. Inoculation was performed on the abaxial surface near the veins of each lobe. Individual drops of the zoospore suspension were applied at puncture sites. Inoculation was carried out before flowering [3]. Infected leaves were stored in airtight glass containers at 20–25°C and periodically moistened. Disease assessment was performed on days 3–4 post-inoculation. The assessors estimated the percentage of infected leaf area [5] and converted it to a 1–9 ordinal severity (resistance) scale where 9 denotes a healthy/most-resistant leaf (minimal infected area) and 1 a heavily infected leaf—i.e., the score is inversely related to infected area. An average score was computed per leaf. Empirically, score 1–2 corresponds to approximately 99% infected area and score 8–9 to approximately 25–29%; Pearson r between score and infected-area fraction is −0.62. The percent-to-score binning (Zhevora et al. 2019) is: 1–3 → 61–100%; 4–5 → 41–60%; 7 → 21–40%; 8–9 → <20%. Photographs were taken and the dataset was compiled on the day of the final assessment.

2.2 Analysed Dataset

The dataset comprised photographs of potato foliage infected with *Phytophthora infestans*. Each photograph contained several leaves; every leaf was split into an individual leaf image, and each leaf comprises up to three upper lobes (leaflets). Analysis rows are at the leaflet level. The processed dataset comprises 78 plants → 140 leaves → approximately 384 leaflets; the Random Forest analysis operates on 138 leaves (after requiring complete features and expert score). The numbers 359, 370, and 350 in the original manuscript were leaflet-level counts at different cleaning stages, which were inconsistently named "images/photographs/leaves"; these have been reconciled to the unit hierarchy above. The images captured the upper three lobes of each potato leaf, with the abaxial surface facing upward against a white background. Each photograph typically contained between three and six leaves. Initial preprocessing consisted of segmenting each photograph into individual leaf images.

2.3 Compiled Dataset for Training a Convolutional Neural Network

The neural network was trained using a dataset of 90 potato leaf images paired with 90 corresponding masks delineating the leaf lobes. This dataset was created by collecting potato leaf images from the Internet and manually crafting masks using an image editor. The masks are white-background images in which each lobe is labeled with one of four colors (red, blue, green, black) such that lobes of the same color do not overlap. The full 90-pair set was not fully retained; 54 pairs remain available. Examples of the created dataset are presented in Figs. 2 and 3.



Figure 2: Example of a source image from the created dataset for leaflet segmentation.



Figure 3: Corresponding mask for the source image, where each color labels a distinct leaflet.

2.4 Developing and Training the CNN for the Leaflet Segmentation

To solve the problem of segmenting potato leaves into individual lobes, we employed a classical encoder-decoder neural network [6]. This architecture gradually compresses complex input data into a low-dimensional latent vector (the encoder), and then reconstructs the output from this representation (the decoder). This model is popular for computer vision tasks, easy to configure for datasets of any size, and was therefore selected as our primary model.

We propose solving the segmentation task by transforming the original three-channel color image into a two-dimensional black-and-white map of isolated leaf lobes, which aligns naturally with the encoder-decoder framework. To increase feature map size in the decoder, we used a simple upsampling layer that duplicates each pixel along the height and width axes. To reduce feature map size in the encoder, we used average pooling, which averages four adjacent pixels to produce a single value [6]. A standard 3×3 convolution was applied at each layer, either halving or doubling the number of channels. Bias terms were omitted to prevent overfitting. The network has 13 convolutional layers with filters $3 \rightarrow 6 \rightarrow 12 \rightarrow 24 \rightarrow 48$ in the encoder, a reshape to a $1 \times 1 \times 768$ bottleneck, and $192 \rightarrow 96 \rightarrow 48 \rightarrow 24 \rightarrow 12 \rightarrow 6 \rightarrow 3 \rightarrow 1$ filters in the decoder, totalling approximately 1.56 million parameters. The proposed network architecture is shown in Fig. 4.

All RGB input images are contrast-equalised using CLAHE prior to resizing to 128×128 and input to the network. This step is applied identically at training, validation and inference. Without CLAHE, segmentation ROC-AUC drops from 0.974 to 0.79.

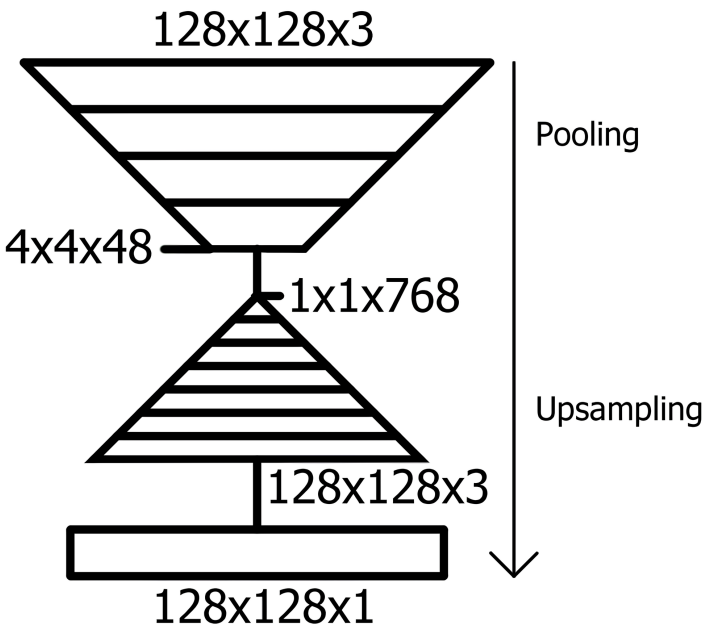


Figure 4: A schematic representation of the developed CNN architecture. The upper part depicts the encoder, and the lower part depicts the decoder.

The upper part of Fig. 4 shows the encoder, where each horizontal strip represents a convolutional block consisting of a convolutional layer, a ReLU activation function [7]: $y = \max(0, x)$, and an average pooling layer. The encoder reduces the input dimensions from $128 \times 128 \times 3$ to $4 \times 4 \times 48$, after which the feature map is flattened into a 768-dimensional vector ($1 \times 1 \times 768$). The lower part of Fig. 4 shows the decoder, where the blocks consist of an upsampling layer (pixel duplication), a convolutional layer, and a ReLU activation function. Additionally, batch normalization layers [8] were applied between all convolutional layers. Training used Adam (lr = 1e−3), binary cross-entropy, and approximately 100 epochs via a 10×10 `fit` loop, with Albumentations augmentations [9] (RGB shift, rotation, elastic transform, channel shuffle, flips, optical distortion, blur). The effective batch size is determined by a custom data generator and is not fixed to a specific value.

2.5 Detecting Spatial Degrees of Phytophthora Infection

A Support Vector Machine (SVM) [10] was used for pixel-wise classification of infection severity. RGB color channels (red, green, blue) were used as features. Each image, represented as a three-dimensional matrix of pixel values across the three channels, was flattened into a set of three-dimensional feature vectors for SVM processing. Like any machine learning algorithm, SVM involves two stages: training and application.

For training mask construction, an expert selected a potato leaf image with late blight symptoms (Fig. 5) from the dataset. This expert-generated mask (Fig. 6) was subsequently used to train the SVM model. The training image contained 78,196 labelled pixels with class balance: class 0: 41%, class 1: 22%, class 2: 28%, class 3: 7%, class 4: 2%. With 5 additional expert-labelled leaves, leave-one-image-out cross-image evaluation gave macro-F1 = 0.35, confirming limited cross-image generalisation (see Section 3.3 and Section Limitations). The color mask was then converted into a one-dimensional array of class labels (integers 0 through 4).



Figure 5: A set of leaves infected by *Phytophthora infestans* used for SVM training.



Figure 6: An expert markup of infection degrees for the SVM training. Different colours correspond to different classes of infection severity.

2.6 Predicting the Expert Score of Samples

To automatically convert the marked-up images into expert-equivalent severity scores, we prepared a training dataset from the collected image collection. The dataset consisted of 138 leaves with complete features and one expert score per leaf. Experts evaluated these images and assigned severity scores on a scale from 1 to 9 [5]. The data were split into training (70%) and test (30%) sets at the leaf level, stratified by rounded expert scores, with random_state = 42.

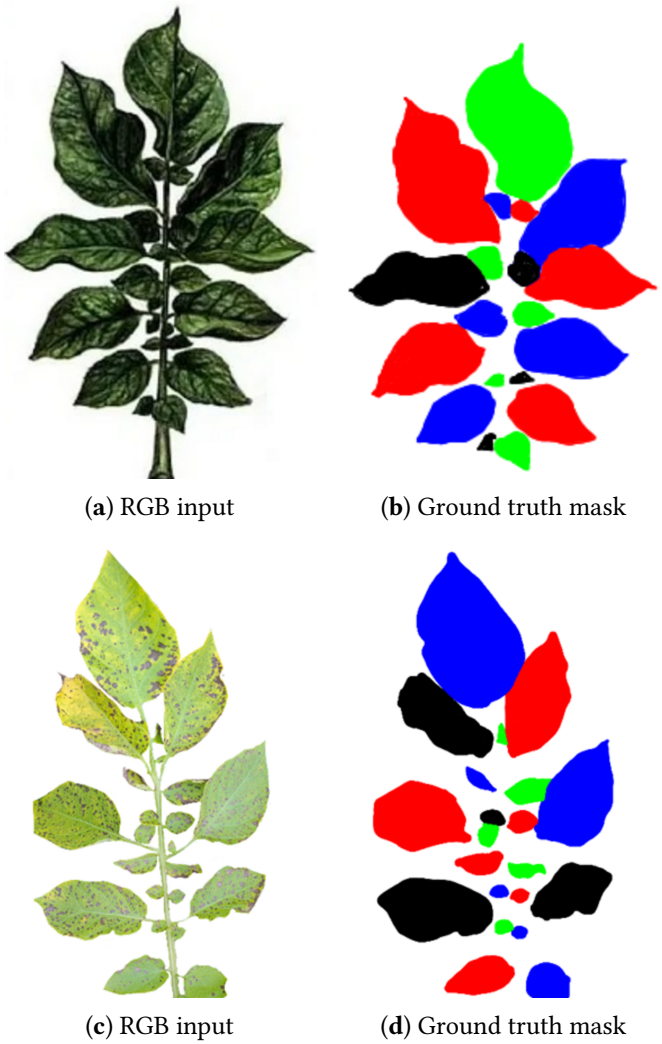
We used a Random Forest classifier [11] for this task. This method is robust to outliers due to its ensemble of uncorrelated decision trees—a critical property when processing data subject to human assessment variability. The input features for the Random Forest were derived from the outputs of the previous pipeline stages (segmentation and infection marking). The final model used 9 colour-area features: green_ratio, yellow_ratio, orange_ratio, brown_ratio, red_ratio, disease_ratio, yo_ratio (yellow+orange), yok_ratio (yellow+orange+brown), and yokr_ratio (yellow+orange+brown+red). Hyperparameters were tuned via cross-validation. The final hyperparameters, confirmed from the training code, are: n_estimators = 10, max_depth = 2, max_features = 1.0 (default), squared-error criterion, random_state = 42. An ablation

study confirmed that cross-validated MAE is essentially constant across 5, 10, 50, 100, and 500 trees (0.71–0.74), so 10 trees is an adequate, parsimonious choice.

3 Results

3.1 Dataset Compilation

To train the convolutional neural network for leaflet segmentation, we created a dataset consisting of 90 image-mask pairs: original potato leaf images and their corresponding manually annotated masks. Eighty pairs were used for training and ten for validation. Fig. 7 shows representative samples from this dataset.



(**a**) RGB input (**b**) Ground truth mask

(**c**) RGB input (**d**) Ground truth mask

Figure 7: Manually annotated dataset for leaflet segmentation. (**a**,**c**) Original RGB images of potato leaves; (**b**,**d**) Corresponding ground truth masks with individual leaflets labeled in distinct colors (red, blue, green, black) for segmentation.

3.2 Segmentation Performance

The CNN model was evaluated using both regression and classification metrics on a pixel-by-pixel basis [12], treating the predicted and ground truth masks as one-dimensional vectors. The results are summarized in Table 1.

In addition, on the 10-pair validation set with CLAHE preprocessing, leaf-class metrics are: IoU = 0.68, Dice = 0.77, and per-leaflet recall = 0.90 (3 of 29 leaflets missed). However, the leaf occupies only approximately 30% of each frame, so overall pixel accuracy (0.92) is partly inflated by background; a trivial all-background predictor already reaches approximately 0.68–0.70 pixel accuracy. IoU/Dice provide a more stringent assessment. The validation set is small (approximately 3–4 independent leaves augmented to 10) and not an independent held-out test; segmentation generalisation is unverified (see Section Limitations).

Visually, the network failed to detect three leaflets among the ten validation images. The network produced no phantom (false) leaflets; all validation masks and their corresponding original images are shown in Fig. 8.

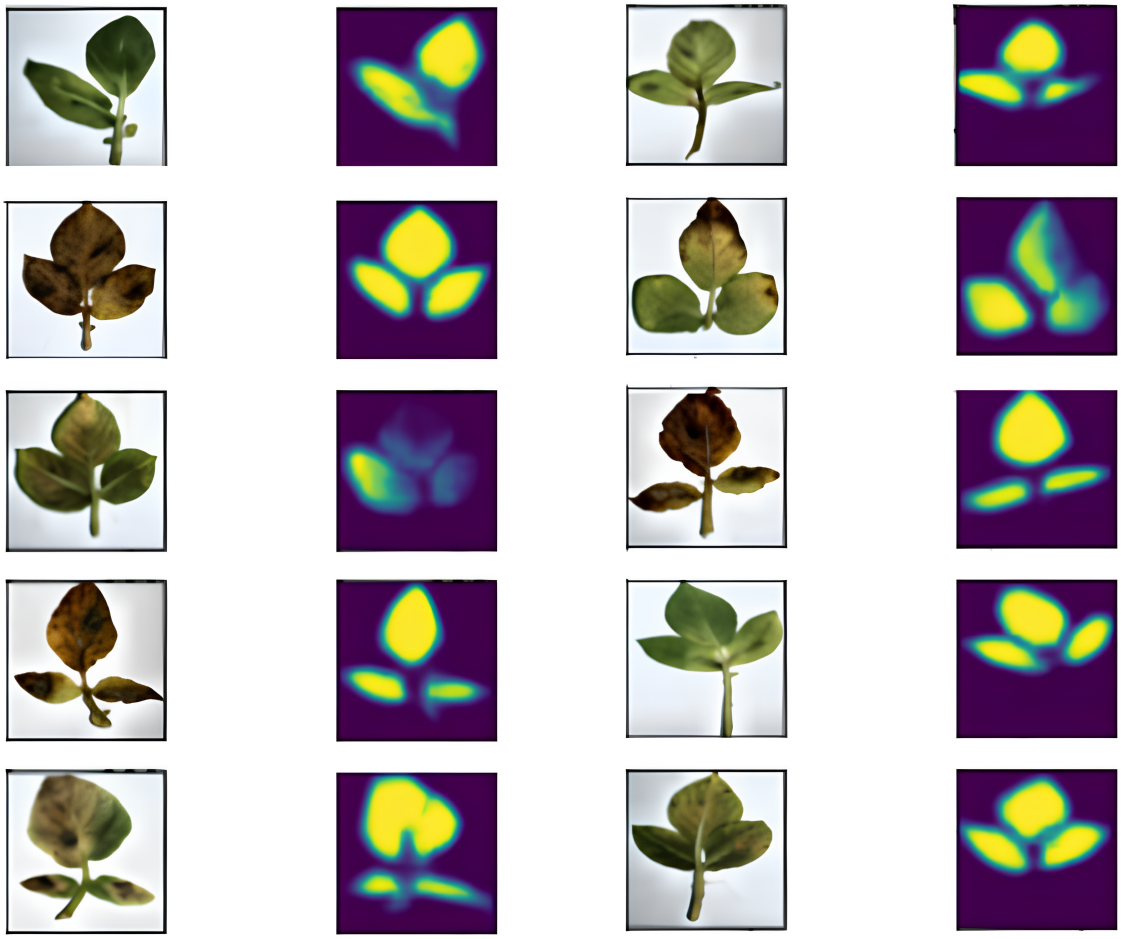


Figure 8: All validation images and their corresponding predicted masks, demonstrating the general quality of the segmentation.

Table 1: Leaflet segmentation performance scores on the validation set.

Metric	Score
Mean Absolute Error	0.076
Mean Squared Error	0.0325
ROC AUC Score	0.974
Accuracy Score	0.92

3.3 Infection Severity Classification

For pixel-wise infection severity classification, 35 potato leaf images showing late blight symptoms were processed. Within-frame cross-validation gives macro-F1 approximately 0.82 and OvR ROC-AUC approximately 0.96; errors concentrate between adjacent severity classes. With 5 additional labelled leaves, leave-one-image-out cross-image evaluation gives macro-F1 = 0.35 (accuracy 0.71). Cross-image generalisation is therefore limited (see Section Limitations). Representative examples are presented in Fig. 9.



(**a**) RGB input (**b**) RGB input (**c**) RGB input

(**d**) Classified output (**e**) Classified output (**f**) Classified output

Figure 9: Examples of SVM-based infection severity classification. (**a**–**c**) Original RGB images of potato leaves inoculated with *Phytophthora infestans*; (**d**–**f**) Corresponding pixel-wise classification maps generated by the trained SVM model. Color coding: green—healthy tissue, yellow—light chlorosis, orange—moderate necrosis, brown—severe necrosis, red—sporulation areas.

3.4 Expert Score Prediction

The Random Forest model was evaluated on a 30% held-out leaf split (stratified by score; random_state = 42; 97 train/41 test of 138 leaves). Both the linear Pearson correlation and the nonlinear Spearman correlation coefficients were computed, along with 95% confidence intervals to assess the statistical reliability of the results. Additional metrics (MAE, RMSE, Bland-Altman) and a 1–9 category confusion matrix were also computed. As shown in Table 2, the updated leaf-level results are: Pearson r = 0.821, Spearman ρ = 0.845, MAE = 0.82, RMSE = 1.11. Bland–Altman analysis shows bias −0.02 and 95% limits of agreement [−1.88, +1.97]. The 1–9 category confusion matrix shows exact match 45.7% and within ±1 = 89.1%. The Spearman value has been updated from 0.78 to 0.845 based on corrected leaf-level validation.

Fig. 10 visualizes the correlation between expert assessments and model predictions on the test set.

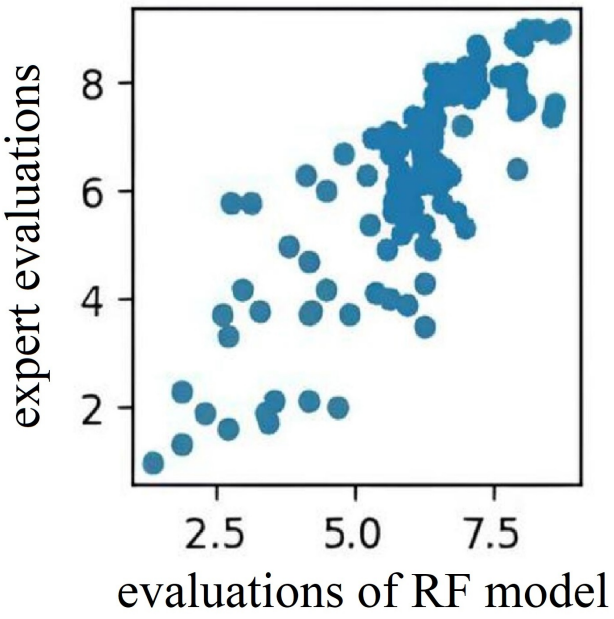


Figure 10: Correlation plot between expert assessments and predicted values from the Random Forest model.

Table 2: Correlation of predicted scores with expert assessments (leaf-level, n = 138).

Correlation	Value	Low CI (95%)	High CI (95%)
Pearson	0.821	0.76	0.87
Spearman	0.845	0.78	0.89

3.5 Complete Processing Pipeline

The proposed algorithm for quantifying infection severity and predicting expert assessments consists of the following steps:

1. Acquire a photograph of at least 1000 × 1000 pixels.
2. Create a downsampled copy of the original image at 128 × 128 pixels using bilinear interpolation.
3. Apply CLAHE contrast equalisation to the downsampled image.
4. Process the CLAHE-equalised image with the trained segmentation CNN to obtain a 128 × 128 pixel black-and-white mask.
5. Resize the mask to the original image dimensions using bilinear interpolation.
6. Apply a brightness threshold of 0.75 to obtain a binary map of leaf lobes.
7. Identify connected components on the binary map to isolate individual leaf lobes, each associated with its own mask region.
8. Apply the trained SVM model to each leaf lobe to generate an infection severity map.
9. For each lobe, convert infection classes to colors for visualization and calculate the proportion of each severity class.
10. Apply the trained Random Forest model to predict the expert severity score.

The approach was implemented in Python as a console application using the TensorFlow/Keras [13], Scikit-learn [14], Scikit-image [15], SciPy [16], NumPy [17], Pandas [18], and Albumentations [9] libraries. The program accepts potato leaf images as input and outputs (i) individual images of each segmented leaf

lobe and (ii) a summary table of infection severity per lobe. The application has been updated to include the CLAHE preprocessing step to match the training/evaluation protocol. An example of the program's output is shown in Fig. 11.

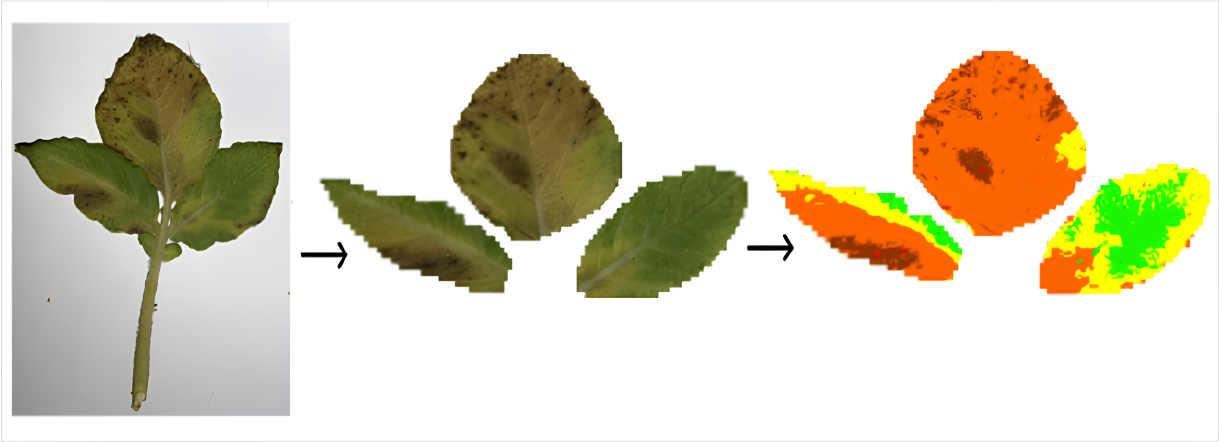


Image	Desease	Leaf	Infection degree	Leaf area	Desease area	Persent of desease
45q3q3	Phytophthora	0	0	105631	109	0.103
45q3q3	Phytophthora	0	1	105631	4236	4.010
45q3q3	Phytophthora	0	2	105631	89951	85.155
45q3q3	Phytophthora	0	3	105631	10671	10.102
45q3q3	Phytophthora	0	4	105631	664	0.628
45q3q3	Phytophthora	1	0	32001	3252	10.162
45q3q3	Phytophthora	1	1	32001	7780	24.311
45q3q3	Phytophthora	1	2	32001	16594	51.854
45q3q3	Phytophthora	1	3	32001	4177	13.052
45q3q3	Phytophthora	1	4	32001	198	0.618
45q3q3	Phytophthora	2	0	50227	18491	36.814
45q3q3	Phytophthora	2	1	50227	24416	48.611
45q3q3	Phytophthora	2	2	50227	7165	14.265
45q3q3	Phytophthora	2	4	50227	155	0.308

Figure 11: Example output of the developed console application.

4 Discussion

This paper presents a multi-stage automated pipeline for assessing late blight severity in potato foliage, integrating three machine learning components: (i) a CNN for segmenting whole leaves into individual leaflets, (ii) an SVM for pixel-wise classification of infection severity, and (iii) a Random Forest for predicting expert severity scores. Additionally, we have released a compiled dataset and a demonstration application.

Several datasets were created for this study. The first, used for SVM training, consists of one photograph of an infected potato leaf and its corresponding expert-annotated mask with five severity classes. The second dataset, comprising 90 image-mask pairs, was designed for CNN training and validation. For

Random Forest training, 138 leaf-level samples processed by the SVM pipeline were used, paired with expert severity scores.

Each component demonstrates sufficient quality metrics for practical application. Our pipeline achieves predictions that correlate strongly with expert assessments (Pearson's r = 0.821; Spearman's ρ = 0.845), supporting its potential use under laboratory conditions. For comparison, Gao et al. [19] report R^2 = 0.655 (visual score vs. lesion count) and lesion IoU = 0.386 on field data; Liu et al. [20] report r ≈ 0.96 vs. experts for wheat stripe rust (continuous severity); Yuan et al. [21] report leaf mAP50 of 97.4% and lesion mIoU of 82.65% for potato late blight severity grading using YOLOv8+UNet3+. Our result (r ≈ 0.82–0.85) is competitive for an ordinal potato resistance scale.

A Random Forest ablation study confirmed that cross-validated MAE is essentially constant across 5, 10, 50, 100, and 500 trees (0.71–0.74)—performance plateaus immediately, so 10 trees is an adequate, parsimonious choice. Feature importance analysis shows the brown (severe-necrosis) area fraction is the dominant predictor (importance approximately 0.56); yellow (0.16), red (0.15), and orange (0.14) contribute substantially less. Cumulative feature groups improved MAE from 1.12 (yellow+orange) to 0.79 (+brown) to 0.73 (+red). A depth ablation on the as-built configuration (10 trees, 9 features) shows max_depth = 2 reproduces the reported correlation (Pearson approximately 0.83 in grouped cross-validation); slightly deeper trees (3–5) reduce MAE modestly (approximately 0.80 to 0.69) before plateauing. The choice of depth = 2 is therefore conservative rather than optimal, and deeper trees represent potential future tuning.

Limitations

Several limitations should be noted. First, all severity scores were assigned by a single expert, so inter-observer agreement cannot be reported. Second, the dataset comprises detached leaves on white background under controlled lighting; independent units are 140 leaves/78 plants (rows are leaflets); external/field validity is untested. Third, SVM cross-image generalisation is limited—leave-one-image-out over 6 labelled leaves gives macro-F1 = 0.35 vs. 0.82 within-frame. Fourth, the CNN validation set is small (approximately 3–4 independent leaves augmented to 10) and not an independent held-out test; segmentation generalisation is unverified. Fifth, the full 88–90 CNN training pairs were not retained (54 pairs survive). Sixth, the CLAHE preprocessing step is essential for the reported CNN performance; the console application has been updated accordingly.

5 Conclusions

In this study, we successfully developed and validated an automated pipeline for the severity quantification and grading of visible *Phytophthora infestans* symptoms in potato leaves. The approach integrates a convolutional neural network for precise leaflet segmentation, a support vector machine for pixel-wise infection grading, and a random forest for predicting expert assessment scores. The results demonstrate high accuracy in spatial infection mapping and a strong correlation between predicted scores and expert evaluations (Pearson's r = 0.821, Spearman's ρ = 0.845), effectively mitigating the subjectivity inherent in traditional visual assessments. This tool offers a robust, reproducible, and efficient method for high-throughput phenotyping, which can significantly reduce human error and labor in potato breeding and phytopathological research. Future work will focus on expanding dataset diversity, improving cross-image generalisation of the SVM component, and adapting the methodology to other crop-disease systems.

Acknowledgement: None.

Funding Statement: This research was funded by the Ministry of Science and Higher Education of the Russian Federation, Agreement 252 No. 075-15-2025-581, dated 23 June 2025.

Author Contributions: The authors confirm contribution to the paper as follows: Conceptualization, Nadezhda Onezhko and Mikhail Divashuk; methodology, Nadezhda Onezhko and Aleksey Ermolaev; software, Aleksey Ermolaev and Daniil Ulyanov; validation, Mariia Samarina, Sergey Avdeev and Gennady Karlov; formal analysis, Nadezhda Onezhko; investigation, Mariia Samarina; resources, Mikhail Divashuk; data curation, Mariia Samarina; writing—original draft preparation, Nadezhda Onezhko; writing—review and editing, Nadezhda Onezhko and Mikhail Divashuk; visualization, Nadezhda Onezhko; supervision, Mikhail Divashuk; project administration, Mikhail Divashuk. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The data that support the findings of this study are available from the Corresponding Author, Nadezhda Onezhko, upon reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: Given his role as Guest Editor of this journal, Mikhail Divashuk had no involvement in the peer review of this article and had no access to information regarding its peer review. Full responsibility for the editorial process for this article was delegated to another journal editor. The authors declare no other conflicts of interest.

References

1. Filipov AV. Fitoftoroz kartofelya [Phytophthorosis of potato]. Moscow, Russia: Vserossiyskiy Nauchno-Issledovatelskiy Institut Fitopatologii (All-Russian Research Institute of Phytopathology, VNIIF); 2012. 88 p. (Prilozhenie k zhurnalu "Zashchita i Karantin Rasteniy" No. 5 [Supplement to the journal "Plant Protection and Quarantine" No. 5]). [cited 2026 Jul 29]. Available from: https://vniif.ru/docs/%D0%9F%D1%83%D0%B1%D0%BB%D0%B8%D0%BA%D0%B0%D1%86%D0%B8%D0%B8%20%D1%81%D0%BE%D1%82%D1%80%D1%83%D0%B4%D0%BD%D0%B8%D0%BA%D0%BE%D0%B2/%D0%9E%D1%82%D0%B4%D0%B5%D0%BB%20%D0%91%D0%9A%D0%B8%D0%9E%D0%9A/%D0%A4%D0%B8%D0%BB%D0%B8%D0%BF%D0%BF%D0%BE%D0%B2_%D0%A4%D0%B8%D1%82%D0%BE%D1%84%D0%BE%D1%82%D0%BE%D1%80%D0%BE%D0%B7%20%D0%BA%D0%B0%D1%80%D1%82%D0%BE%D1%84%D0%B5%D0%BB%D1%8F.pdf. (In Russian).
2. Hannibal FB, Gasich EL, Orina AS. Otsenka ustoychivosti selektsionnogo materiala krestotsvetnykh i paslenovykh kul'tur k al'ternariozam [Assessment of Resistance of Breeding Material of Cruciferous and Solanaceous Crops to Alternaria Diseases]. St. Petersburg, Russia: VIZR; 2011. 40 p. (Methodical manual). [cited 2026 Jul 29]. Available from: http://vizrspb.ru/assets/images/fitopatology_lab/documents/%D0%9C%D0%B5%D1%82%D0%BE%D0%B4.%D0%BF%D0%BE%D1%81%D0%BE%D0%B1%D0%B8%D0%B5%20-%20%D0%BE%D1%86%D0%B5%D0%BD%D0%BA%D0%B0%20%D1%83%D1%81%D1%82%D0%BE%D0%B9%D1%87%D0%B8%D0%B2%D0%BE%D1%81%D1%82%D0%B8%202011.pdf. (In Russian).
3. Zhevora SV, Fedotova LS, Starovoitov VI, Zeyruk VN, Korshunov AV, Pshechenkov KA, et al. Metodika provedeniya agrotekhnicheskikh opytov, uchetov, nablyudeniy i analizov na kartofele [Methods of Agrotechnical Experiments, Surveys, Observations and Analyses on Potatoes]. Moscow, Russia: FGUP "Nauka" Publishing House; 2019. 120 p. [cited 2026 Jul 29]. Available from: https://elibrary.ru/item.asp?id=41770505. (In Russian).
4. DeChant C, Wiesner-Hanks T, Chen S, Stewart EL, Yosinski J, Gore MA, et al. Automated identification of northern leaf blight-infected maize plants from field imagery using deep learning. Phytopathology. 2017;107(11):1426–32. [CrossRef]
5. Sidorik AI, Khasanov VT. Otsenka porazhennosti virusnymi zabolevaniyami i polevoy ustoychivosti k fitoftorozu sortoobraztsov kartofelya kollektsionnogo pitomnika Kostanayskogo NIISKh [Evaluation of viral diseases and field resistance to late blight in potato accessions from the collection nursery of Kostanai Agricultural Research Institute]. Vestnik Nauki Kazakhskogo Agrotekhnicheskogo Universiteta imeni S. Seifullina. 2017;4(95):42–51. [cited 2026 Aug 11]. Available from: https://rmebrk.kz/journals/4561/48004.pdf#page=42. (In Russian).

6. Simard PY, Steinkraus D, Platt JC. Best practices for convolutional neural networks applied to visual document analysis. In: Proceedings of the 7th International Conference on Document Analysis and Recognition; 2003 Aug 3–6; Edinburgh, UK. [CrossRef]
7. Rössig A, Petkovic M. Advances in verification of ReLU neural networks. J Glob Optim. 2021;81(1):109–52. [CrossRef]
8. Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: Proceedings of the 32nd International Conference on Machine Learning; 2015 Jul 6–11; Lille, France. p. 448–56. [CrossRef]
9. Albumentations Library [Internet]. [cited 2024 Jun 30]. Available from: https://albumentations.ai/docs/.
10. Cervantes J, Garcia-Lamont F, Rodríguez-Mazahua L, Lopez A. A comprehensive survey on support vector machine classification: Applications, challenges and trends. Neurocomputing. 2020;408:189–215. [CrossRef]
11. Breiman L. Random forests. Mach Learn. 2001;45:5–32. [CrossRef]
12. Boutaba R, Salahuddin MA, Limam N, Ayoubi S, Shahriar N, Estrada-Solano F, et al. A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities. J Internet Serv Appl. 2018;9(1):1–99. [CrossRef]
13. TensorFlow Library [Internet]. [cited 2024 Jun 30]. Available from: https://www.tensorflow.org/.
14. Scikit-learn Library [Internet]. [cited 2024 Jun 30]. Available from: https://scikit-learn.org.
15. Scikit-image Library [Internet]. [cited 2024 Jun 30]. Available from: https://scikit-image.org/docs/stable/.
16. Scipy Library [Internet]. [cited 2024 Jun 30]. Available from: https://docs.scipy.org/doc/scipy/.
17. NumPy Library [Internet]. [cited 2024 Jun 30]. Available from: https://numpy.org/doc/.
18. Pandas Library [Internet]. [cited 2024 Jun 30]. Available from: https://pandas.pydata.org/docs/.
19. Gao J, Westergaard JC, Sundmark EHR, Bagge M, Liljeroth E, Alexandersson E. Automatic late blight lesion recognition and severity quantification based on field imagery of diverse potato genotypes by deep learning. Knowl-Based Syst. 2021;214:106723. [CrossRef]
20. Liu W, Chen Y, Lu Z, Lu X, Wu Z, Zheng Z, et al. StripeRust-Pocket: A two-stage deep learning pipeline for wheat stripe rust severity assessment. Plant Phenomics. 2024;6:201. [CrossRef]
21. Yuan P, Jiang L, Cheng Z, Tan Y, Yang Y, He C. Potato late blight severity assessment using YOLOv8 and UNet3+. Front Plant Sci. 2025;16:1616864. [CrossRef]