

Smart Farming with Deep Learning: Effective Soybean Leaf Disease Diagnosis via InceptionV3

Salman Ahmad

Author's Affiliation

King Saud University, Riyadh, Saudi Arabia

*Corresponding Author:

Salman Ahmad

King Saud University, Riyadh, Saudi Arabia.

Keywords:

Soybean, Machine Learning, CNN, InceptionV3, Caterpillar, *Diabrotica speciosa*

Received on 13.06.2026

Received on 21.08.2026

Accepted on 09.09.2026

Abstract

Soybean is a significant crop worldwide due to its high protein and edible oil content. However, foliar injury caused by insect pests, such as caterpillars and *Diabrotica speciosa*, can reduce its productivity. Quick and reliable detection of such damage can help make timely crop protection decisions. This work investigates the use of a convolutional neural network (CNN) based on a pre-trained InceptionV3 model for identifying soybean leaf conditions. Transfer-learning was used with InceptionV3 pre-trained on ImageNet. The model was modified to identify three categories of soybean leaves: Healthy, Caterpillar, and *Diabrotica speciosa*. The study is based on the Mendeley dataset, which consists of 6,410 images. The InceptionV3 feature extractor, Global Average Pooling, Dense and Dropout layers were added for classification. The images available were split into 80% training and 20% validation. Standard classification measures were used to analyze model performance. The modified InceptionV3 network achieved good classification results for the soybean-leaf dataset. The study achieves an overall test accuracy of around 94%, and high precision, recall and F1-score values for the three classes. The results show that an ImageNet-pre-trained architecture can be adapted to this crop-specific image-recognition task, and could be useful for technology-assisted crop monitoring.

How to cite this article: Ahmad S. (2026). Smart Farming with Deep Learning: Effective Soybean Leaf Disease Diagnosis via InceptionV3. *Bulletin of Pure and Applied Sciences-Botany*, 45B (2), 31-40.

1. INTRODUCTION

In recent years, machine learning (ML) and computer-vision (CV) methods have been applied to solve agricultural issues that are hard to efficiently solve by manual observation (Najdenovska et al., 2021). Legume crops are especially significant due to their role in providing protein, fibre, and other nutritious components to human diets, as well as their ability to enhance soil fertility through biological nitrogen fixation (BNF) (Jat et al., 2024). Although these advantages exist, pathogens and insect pests can impact crop performance. Soybean is a protein and oil source crop that is susceptible to multiple leaf-feeding pests, and symptoms may include discoloration, necrotic

areas, spots, and defoliation (He et al., 2020; Kim et al., 2024). Fungal pathogens are also responsible for crop-health issues, and recently, plant-derived antifungal activity has been investigated against *Fusarium oxysporum* f. *lycopersici*, a pathogen of agricultural relevance (Beg et al., 2025). *Spodoptera litura* has also been studied as an economically significant leaf-feeding pest, highlighting the importance of developing better pest management strategies (Sharma & Kulkarni, 2024). Early detection of damage can thus help to make more timely management decisions.

Traditional methods of leaf damage diagnosis may rely on visual inspection or laboratory testing, which can be time-consuming and labor-

intensive. An alternative is offered by AI and deep learning, which can automatically identify image-based symptoms (Kaur & Singh, 2022). CNNs are particularly useful because the relevant visual features can be learned directly from image data, without being designed completely by hand. Previous image processing methods were more dependent on enhancement, segmentation, filtering and manually selected descriptors (Barbedo, 2013; Al-Hiary et al., 2011). However, in deep-learning applications, performance is still dependent on the quality of the dataset, the method of annotation, the validation strategy, and the avoidance of information leakage between data subsets, which have also been highlighted in image-classification research outside of agriculture (Taqvi, 2025).

CNN development has progressively improved automatic visual feature learning. AlexNet demonstrated the effectiveness of deep convolutional architectures for large-scale image recognition (Krizhevsky et al., 2012), while later networks such as GoogLeNet/Inception (Szegedy et al., 2015), DenseNet (Huang et al., 2017), and related architectures introduced increasingly efficient feature-reuse and multi-scale processing strategies. InceptionV3 is particularly attractive for image classification because its architecture combines factorized convolutions and parallel operations that can represent visual information at multiple scales. Such characteristics are relevant to soybean leaves, where pest damage may differ only in small local texture or shape patterns. Transfer learning further allows these learned representations to be adapted to agricultural datasets without training an entire network from the beginning. Consistent with this approach, a recent study of rice-leaf disease recognition used a fine-tuned EfficientNetB0 network with data augmentation and reported 98.2% accuracy (Tholupuluri et al., 2024).

Deep-learning approaches have already been explored for several plant-disease and crop-monitoring problems. Walleign et al. (2018), for instance, applied a CNN to soybean disease images from PlantVillage and reported high classification performance. Mohanty et al. (2016) also showed that transfer learning can be effective for image-based plant-disease

recognition. More recent work has moved toward attention-enhanced and modern backbone architectures. Wu et al. (2023) proposed an improved ConvNeXt-based method for soybean leaf-disease classification, while Goshika et al. (2024) applied a deep-learning model to classify and assess soybean leaf damage. These studies illustrate the continuing shift from conventional feature engineering toward trainable models capable of learning fine-grained visual differences. Jung et al. (2023) further demonstrated the use of CNN-based classification with local interpretable model-agnostic explanation methods in plant-species discrimination, highlighting the growing interest in combining prediction with interpretable image analysis.

Relative to traditional image-processing pipelines, deep-learning systems can provide a more flexible framework for adapting to complex image datasets. Cho (2024) compared machinelearning methods for agricultural leaf-disease recognition, and recent crop research is still investigating transfer learning, which involves fine-tuning a pre-trained network with less task-specific training than training a network from scratch. However, in field conditions, uneven illumination, background clutter, seasonal changes, and variation in image acquisition can affect performance. Agricultural systems can also be affected by broader environmental conditions, such as rainfall variability and drought-related conditions, which have direct implications for crop production and resource availability (Golekar & Diwate, 2024). The utility of CNN-based crop analysis is that it can help in more rapid monitoring and more targeted interventions. Automated recognition tools can be integrated into digital or precision-agriculture systems to assist farmers in detecting damaged plants at an earlier stage and optimizing input use. The integration of sensing, data processing, and decision support is becoming more common in precision and smart farming, aiming to enhance resource management and productivity (Geetanjali & Singh, 2025). Such systems can be further expanded by mobile, cloud and edge-computing platforms. Further advances will rely on larger and more diverse datasets, more robust field testing of validation, and ongoing

testing of attention mechanisms and other architectures for challenging real-world imagery.

2. MATERIALS AND METHODS

2.1 Dataset

This work uses an image collection where the leaves of soybeans are labeled as Caterpillar, *Diabrotica speciosa*, and Healthy. The images were obtained from the Mendeley Data

repository described by Mignoni (2021). Three representative examples are shown in Figure 1 for each of the three categories. The photographs were taken in natural light and may show variations in lighting, shadows, and other environmental factors. The total number of images in the dataset is 6410. Of these, 3,309 are in the Caterpillar class, 2,205 are in the *Diabrotica speciosa* class, and 896 are in the Healthy class.




Figure 1: Representative soybean leaf images: (a) Healthy, (b) Caterpillar-damaged, and (c) *Diabrotica speciosa*-affected.

The images were captured under realistic conditions, resulting in meaningful variation in background, illumination, leaf appearance, and pest-related symptoms. This variability is crucial if a model is required to distinguish between similar categories, instead of learning a limited number of image patterns. The soybean phyllosphere is also biologically complex, and contains a variety of microbial communities that are affected by biotic and abiotic factors (Arora et al., 2025), further demonstrating the natural variability of leaf surfaces.

2.2 Data preprocessing

The image preparation process was divided into three main steps before the model training:

Resizing and Rescaling: The original images were not of the same spatial resolution. They were then resized to 224×224 pixels to ensure that each sample was the same size as the input to the model. The pixel values were also normalized to ensure that the input data was on a consistent numerical scale and to ensure stable optimization during training.

Data augmentation: Extra training variability was added by rotating, flipping horizontally or vertically, zooming and changing brightness. These transformations can reveal the model to several plausible versions of the same visual pattern, and can help to prevent overfitting by discouraging the model from relying on a single orientation or lighting condition.

Dataset splitting: The image collection was split into 80% training and 20% validation. This setup enabled the majority of images to be used for parameter learning and a separate subset for validation. The Results section also presents performance on a test set of 641 images, but the source manuscript does not further explain how this test set was obtained from the split mentioned above.

2.3 Model Selection

InceptionV3, initialized with ImageNet weights, was chosen as the base network. Its hierarchical and multi-scale feature extraction is well suited to a task in which the distinction between a healthy leaf and pest-related damage may

depend on relatively small local visual differences.

2.4 Model Architecture

The classification system uses InceptionV3 as a pre-trained feature extractor and removes the original top classifier so that the network can be adapted to three soybean-leaf categories (Chollet, 2017). Inception blocks process image features through parallel convolutional paths, including operations based on different receptive-field sizes. Factorized and asymmetric convolutions are used within the architecture to lower computational cost while preserving expressive feature extraction. The adapted network therefore retains general visual representations learned from ImageNet while adding task-specific layers for the soybean dataset. After the InceptionV3 backbone, task-specific layers were added to convert the extracted feature maps into class predictions.

Global Average Pooling was used to summarize each final feature map into a compact representation before classification. This

strategy, introduced in the Network in Network framework by Lin et al. (2014), reduces the need for large fully connected layers and links learned feature maps more directly to class-level evidence. In the present architecture, the pooled features are passed to a dense layer containing 512 units with ReLU activation so that task-specific combinations of the extracted features can be learned.

A Dropout layer with a rate of 0.5 was then applied as a regularization mechanism. Dropout reduces co-adaptation by randomly suppressing a proportion of activations during training, which can improve generalization (Hinton et al., 2012). The last Dense layer has three output units with softmax activation to represent the three classes of soybean-leaf. The model description also contains batch-normalization and auxiliary-classifier concepts that are designed to help with stable optimization and to prevent overfitting. The overall architecture is shown in Figure 2, and the reported layer outputs and number of parameters are summarized in Table 1.




Figure 2: InceptionV3-based model architecture with additional Global Average Pooling, Dense, and Dropout layers.

Table 1: Parameters used in the InceptionV3-based model structure.

Layer Type	Output Shape	Parameters
InceptionV3	(None, 1, 1, 2048)	21,802,784
GlobalAveragePooling2D	(None, 2048)	0
Dense	(None, 512)	1,049,088
Dropout	(None, 512)	0
Dense (Output Layer)	(None, 3)	1,539

2.5 Evaluation parameters

Evaluation parameters several complementary classification measures were used to evaluate performance, not just accuracy. The analysis was based on the confusion matrix, Matthews Correlation Coefficient (MCC), accuracy, precision, recall, and F1-score, following the common practice in machine learning (Rainio et al., 2024). These metrics collectively describe the overall correctness and the prediction behaviour per class. 1.

1. Confusion matrix

A confusion matrix is a table that displays the relationship between the predicted labels and the actual labels. It can be summarized by the number of True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN). In this study, the matrix allows the identification of the classes of the species correctly identified and the classes that are confused between the classes of the species Healthy, Caterpillar and *Diabrotica speciosa*.

		Predicted Values	
		Positive	Negative
Actual Values	Positive	TP	FN
	Negative	FP	TN

2. Matthews Correlation Coefficient (MCC)

MCC offers a correlation-based summary of agreement between observed and predicted labels. It can be either negative or positive, with +1 being a perfect match, 0 being no correlation, and -1 being a complete disagreement. MCC is informative when class frequencies are unequal, since it uses information from the various outcomes of the confusion matrix.

$$MCC = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

3. Accuracy

Accuracy is the fraction of all evaluated samples that receive the correct class label.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

4. Precision

Precision expresses how many samples predicted as belonging to a given class are actually members of that class and is commonly calculated as $TP / (TP + FP)$.

$$Precision = \frac{TP}{TP + FP}$$

5. Recall

Recall describes the proportion of actual positive samples for a class that the model successfully retrieves and is commonly calculated as $TP / (TP + FN)$. A larger recall value means that fewer true instances of that class are overlooked.

$$Recall = \frac{TP}{TP + FN}$$

6. F1 score

The F1-score combines precision and recall through their harmonic mean, providing a single summary when both false-positive and false-negative errors are important.

$$F1 - score = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

2.6 Training Process

The training was done using the Adam optimizer and categorical cross-entropy loss, which is appropriate for the three-class prediction problem. The network was trained for 25 epochs. To track changes in learning behaviour and potential overfitting, accuracy and loss were tracked for both the training and validation subsets during training.

3. RESULTS AND DISCUSSION

This section summarizes the training, validation, and test results reported. The accuracy and loss curves are shown in Figure 3(a, b) over 25 epochs. The training accuracy was reported as 98.45% and the validation accuracy was 95.79% at the end of training. The test set is claimed to be 94.70% accurate. These values show that the model has good predictive ability when tested with samples other than those used directly for parameter fitting.




Figure 3: The accuracy and loss after 25 epochs for training and validation are shown in Figure 3(a) and 3(b), respectively.

The confusion matrix for the three soybean-leaf categories is shown in Figure 4. The majority of the observations are on the main diagonal, indicating that most of the samples were correctly labeled. The smaller off-diagonal counts represent errors across classes, and show that misclassification was relatively small

compared to the number of correct predictions. The reported MCC value is about 0.91. This value is a measure of the correspondence between the predicted and actual labels for the evaluated classes on the -1 to +1 MCC scale, and indicates a high level of correspondence.




Figure 4: Confusion matrix for soybean leaf-condition classification.

The reported MCC value is approximately 0.91. On the -1 to +1 MCC scale, this value reflects a high level of correspondence between the

predicted and actual labels across the evaluated classes.

Table 2: Classification precision, recall, F1-score, and support

Classification	Precision	Recall	F1-score	Support
Caterpillar	0.96	0.94	0.95	331
<i>Diabrotica speciosa</i>	0.93	0.94	0.93	220
Healthy	0.90	0.95	0.92	90
Accuracy			0.94	641
Macro Average	0.93	0.94	0.94	641
Weighted Average	0.94	0.94	0.94	641

Table 2 shows the precision, recall, F1-score and support for each class in the Healthy, Caterpillar and *Diabrotica speciosa* classes. The precision values are between 0.90 and 0.96, while the recall values are between 0.94 and 0.95. The F1-scores are between 0.92 and 0.95. The overall reported accuracy is 0.94 for 641 test samples,

while both the macro-average and weighted-average F1-scores are 0.94. The results demonstrate that the model achieved a relatively balanced predictive performance despite the unequal number of images in the three classes.




Figure 5: Receiver Operating Characteristic (ROC) curves of the InceptionV3 model for soybean leaf condition classification

The ROC curves for each class are displayed in Figure 5. The curves show the true-positive rate (TPR) as a function of the false-positive rate (FPR) for different decision thresholds, and the AUC quantifies the discrimination performance. The reported AUC is 0.98 for class 0 and 0.99 for both class 1 and class 2. Values near 1 represent good separation between each target class and the alternatives in the evaluated data set. Previous plant-disease research has employed various CNN architectures, including relatively shallow networks, deeper residual networks,

dense networks, and inception-style networks. Reported performance varies with image quality, class complexity, dataset composition, and experimental design. In the current work, the test accuracy of InceptionV3 was around 94% with the help of transfer learning and multi-scale feature extraction. The main benefit of this setup is that the general visual representations learned from ImageNet can be transferred to soybean images, reducing the number of task-specific training images needed

while retaining fine visual patterns related to leaf damage.

4. CONCLUSION

The use of deep learning and computer vision has expanded the toolbox for image-based crop damage recognition. Automated leaf analysis could be used in conjunction with manual inspection for soybean production, as it can be used to screen a large number of images quickly and at a larger scale. In the present work, the authors use transfer learning to solve a three-class soybean-leaf problem and show how a pre-trained CNN can be fine-tuned to classify healthy leaves from two types of pest-induced damage. The choice of InceptionV3 was made due to its architecture that can represent visual information at multiple spatial scales while maintaining computational efficiency. After adaptation to the soybean dataset, the study reports approximately 94% overall test accuracy. To simplify the model and facilitate generalization, Global Average Pooling and Dropout were added to the task-specific classification head. The resulting system thus consists of a pre-trained feature extractor and a relatively small classifier for the three target classes. There are a number of restrictions that need to be addressed before the approach can be considered a general field-deployment solution. The cost of creating large annotated datasets is high, and performance can vary if lighting, backgrounds, seasons, cameras, or crop conditions vary from those in the training images. For practical use, more diverse data, external validation, and periodic retraining would therefore be important. On a larger scale, AI can be integrated with remote-sensing technologies to enhance automated mapping, environmental monitoring, and data-driven agricultural decision making. Future research may involve expanding the classifier to other soybean diseases and pest-damage types, comparing the current network with attention-based and hybrid networks, and evaluating domain-adaptation techniques for images taken in new environments. Validation in the field, across locations and seasons, would also help to establish the transferability of the reported performance beyond the current dataset. This type of research may help to reinforce the importance of computer vision in precision-

agriculture systems for timely crop-health monitoring and better management decisions.

REFERENCES

- Al-Hiary, H., Bani-Ahmad, S., Reyalat, M., Braik, M., & ALRahamneh, Z. (2011). Fast and accurate detection and classification of plant diseases. *International Journal of Computer Applications*, 17(1), 31-38. <https://doi.org/10.5120/2183-2754>
- Arora, V., Kanwal, A., Sharma, I., & Singh, R. (2025). The phyllosphere and the phylloplane microbial biome: A diverse and dynamic community. *Bio-Science Research Bulletin*, 41(1), 62-69. <https://doi.org/10.5281/zenodo.16157091>
- Barbedo, J. G. A. (2013). Digital image processing techniques for detecting, quantifying and classifying plant diseases. *SpringerPlus*, 2, 660. <https://doi.org/10.1186/2193-1801-2-660>
- Beg, M. A., Ziya, R., Mansoor, F., Rahman, S., Imtiyaz, M., & Beg, M. J. (2025). In vitro antifungal activity of weed plants of Eastern Uttar Pradesh against *Fusarium oxysporum* f. *lycopersici* (Sacc.) Snyder & Hansen. *Bulletin of Pure and Applied Sciences-Botany*, 44B(1), 23-29. <https://bpasjournals.com/botany/index.php/journal/article/view/137>
- Cho, O.-H. (2024). An evaluation of various machine learning approaches for detecting leaf diseases in agriculture. *Legume Research*, 47(4), 619-627. <https://doi.org/10.18805/LRF-787>
- Chollet, F. (2017). *Deep Learning with Python*. Manning Publications.
- Geetanjali, & Singh, A. R. (2025). Precision and smart farming: An approach to efficient resource management in agriculture. *Journal of Artificial Intelligence in Agriculture and Applied Research*, 1(1), 25-33. <https://agrpjournals.com/journals/journal/index.php/jaiaar/article/view/6>

- Golekar, R. B., & Diwate, P. (2024). Assessment of rainfall in monsoon season 2023: A case study of Maharashtra State. *Bulletin of Pure and Applied Sciences-Geology*, 43F(2), 64-82. <https://bpasjournals.com/geology/index.php/journal/article/view/159>
- Goshika, S., Meksem, K., Ahmed, K. R., & Lakhssassi, N. (2024). Deep learning model for classifying and evaluating soybean leaf disease damage. *International Journal of Molecular Sciences*, 25(1), 106. <https://doi.org/10.3390/ijms25010106>
- He, Y., Zhou, Z., Tian, L., Liu, Y., & Luo, X. (2020). Brown rice planthopper (*Nilaparvata lugens* Stal) detection based on deep learning. *Precision Agriculture*, 21, 1385-1402. <https://doi.org/10.1007/s11119-020-09726-2>
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. *arXiv*. <https://doi.org/10.48550/arXiv.1207.0580>
- Huang, G., Liu, Z., van der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4700-4708. <https://doi.org/10.1109/CVPR.2017.243>
- Jat, B. L., Sharma, H. C., Pagaria, P., Meena, A. K., Mali, G. R., & Khan, T. (2024). Legumes: Source of bioactive compounds and their potential use in legume crops improvement: A review. *Legume Research*, 47(11), 1827-1834. <https://doi.org/10.18805/LR-4945>
- Jung, D.-H., Kim, H.-Y., Won, J. H., & Park, S. H. (2023). Development of a classification model for *Cynanchum wilfordii* and *Cynanchum auriculatum* using convolutional neural network and local interpretable model-agnostic explanation technology. *Frontiers in Plant Science*, 14, 1169709. <https://doi.org/10.3389/fpls.2023.1169709>
- Kaur, R., & Singh, S. (2022). A comprehensive review of object detection with deep learning. *Digital Signal Processing*, 132, 103812. <https://doi.org/10.1016/j.dsp.2022.103812>
- Kim, B.-H., Alamri, A. M., & AlQahtani, S. A. (2024). Leveraging machine learning for early detection of soybean crop pests. *Legume Research*, 47(6), 1023-1031. <https://doi.org/10.18805/LRF-794>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097-1105.
- Lin, M., Chen, Q., & Yan, S. (2014). Network in Network. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1312.4400>
- Mignoni, M. E. (2021). Images of soybean leaves [Dataset]. *Mendeley Data*, V1. <https://doi.org/10.17632/bycbh73438.1>
- Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, 7, 1419. <https://doi.org/10.3389/fpls.2016.01419>
- Najdenovska, E., Dutoit, F., Tran, D., Plummer, C., Wallbridge, N., Camps, C., & Raileanu, L. E. (2021). Classification of plant electrophysiology signals for detection of spider mites infestation in tomatoes. *Applied Sciences*, 11(4), 1414. <https://doi.org/10.3390/app11041414>
- Rainio, O., Teuho, J., & Klén, R. (2024). Evaluation metrics and statistical tests for machine learning. *Scientific Reports*, 14, 6086. <https://doi.org/10.1038/s41598-024-56706-x>

- Sharma, A., & Kulkarni, P. (2024). Biogenic synthesis of zinc nanoparticles with indigenous *Penicillium citrinum* IBA5VYT and their effectiveness against cotton leafworm. *Bulletin of Pure and Applied Sciences-Zoology*, 43B(1S), 920-931.
<https://bpasjournals.com/zoology/index.php/journal/article/view/176>
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1-9.
<https://doi.org/10.1109/CVPR.2015.7298594>
- Taqvi, S. (2025). Deep learning for detection and classification of oral mucosal lesions. *Journal of Dentistry and Clinical Sciences*, 1(1), 23-30.
<https://agrpjournals.com/journals/journal/index.php/jdacs/article/view/11>
- Tholupuluri, G. B., Charith, B., & Gopal Krishna, U. M. (2024). Deep learning algorithms for rice leaf disease detection. *Bulletin of Pure and Applied Sciences-Zoology*, 43A(2S), 331-343.
<https://bpasjournals.com/zoology/index.php/journal/article/view/243>
- Walleign, S., Polceanu, M., & Buche, C. (2018). Soybean plant disease identification using convolutional neural network. *Proceedings of the 31st International Florida Artificial Intelligence Research Society Conference (FLAIRS-31)*, 146-151.
- Wu, Q., Ma, X., Liu, H., Bi, C., Yu, H., Liang, M., Zhang, J., Li, Q., Tang, Y., & Ye, G. (2023). A classification method for soybean leaf diseases based on an improved ConvNeXt model. *Scientific Reports*, 13, 19141.
<https://doi.org/10.1038/s41598-023-46492-3>