Leveraging ResNet-50 for Precision Toxicity Classification in Plants: A Vision-Based Approach to Safeguard Public Health

Aria Makhija[1]

[1] Morris County School of Technology, New Jersey, USA
aria.makhija@mcvts.org

Abstract. The classification of toxic and non-toxic plants plays an important role in ensuring public safety, especially in agriculture, food safety, and health. Correct identification of these plants can prevent accidental poisoning and promote ecological protection. In this paper, we investigate the application of the ResNet-50 model for the classification of toxic and non-toxic plants. Leveraging the powerful feature extraction techniques of the ResNet-50 architecture, the model achieved 89.6% accuracy, 87.4% precision, 91.1% recall, and an 89.2% F1 score, demonstrating the model's effectiveness. Transfer learning proved effective with limited data while maintaining high performance metrics in the classification task. Future research could focus on expanding the dataset to include more plant species and exploring other state-of-the-art models to improve classification accuracy. Additionally, integrating these models with mobile applications or monitoring systems could provide solutions for business and public use, enhancing environmental protection and public safety.

Keywords: Poisonous plants, Non-poisonous plants, ResNet-50, Plant classification, Deep learning, Transfer learning, Accuracy.

1 Introduction

It is essential to classify plants as toxic or non-toxic for public health and environmental protection. Toxic plants pose significant risks, especially in agriculture, food safety, and foraging, where accidental ingestion or exposure may sometimes result in severe health complications or death. Accurate identification of plants is thus crucial to avoid poisoning incidents. With the advancements in computer vision and deep learning, it is now possible to automate plant identification in great detail. Direct classification of toxic versus non-toxic plants has more than just academic interest; it contributes directly to public safety and environmental stewardship. For instance, the mistaken identification of toxic plants constitutes a severe public health risk, as exposure rates are high due to plant foraging practiced in rural regions. Additionally, unchecked propagation of toxic species may damage agricultural or ecological systems by poisoning animals or altering indigenous ecosystems (Wendt et al., 2022).

Recent advancements in machine learning and deep neural networks have enabled the automation of plant classification with unprecedented accuracy. This is beneficial not only for professionals but also for the general public (Noor et al., 2022). Among deep learning models, Convolutional Neural Networks (CNNs) have proven very promising for image classification, and transfer learning using pre-trained models such as ResNet-50 has emerged as the most efficient way to obtain reliable results for plant classification.

The deep residual network, ResNet-50, is widely used because it helps reduce the problem of vanishing gradients and extracts hierarchical features from images. This enables the model to have a deep architecture without causing gradient degradation, making it one of the most powerful models for plant classification, where small visual differences are important. This model adapts well to tasks with limited labeled data, utilizing the technique of transfer learning, where a model pre-trained on a large dataset such as ImageNet is fine-tuned for specific classifications. The application of ResNet-50 in classifying poisonous and non-poisonous plants allows for excellent feature extraction from plant images, capturing complex visual details such as leaf shape, texture, and color patterns that are critical for distinguishing between species (Zuhri et al., 2022). This computationally efficient and training-time expedient approach also increases classification accuracy through the use of pre-trained knowledge from a more comprehensive dataset.

Using ResNet-50 provides several advantages for this task. For one, its deeper architecture allows the model to capture and learn complex patterns that traditional approaches may struggle to identify. An example of this is plant classification based on leaf texture and patterns, which has traditionally been challenging due to variations in lighting and seasons (Hassoon and Hantoosh, 2023). The generalization capability of ResNet-50 in handling such variations has made it a preferred choice for plant identification tasks, particularly in differentiating between poisonous and non-poisonous species. Furthermore, this model can efficiently handle large datasets with fewer parameters than its deeper counterparts, ensuring an optimal balance between computational cost and performance (Mezzasalma et al., 2017).

2 Literature Survey

Several studies have used machine learning and deep learning techniques to address the problem of plant classification. H et al. (2024) reviewed the use of machine learning to detect poisonous plants, as such models are capable of identifying minute differences that the human eye may not be able to spot. Similarly, Azadnia et al. (2024) proposed the use of computer vision to identify medicinal and toxic plants based on leaf characteristics, using a deep neural network to improve classification accuracy. These studies highlight the necessity of using advanced neural architectures like ResNet-50 to handle the complexities involved in plant classification tasks.

The ResNet-50 model has gained widespread usage for image classification tasks, including the classification of plants. Residual learning allows the network to learn even more complex representations without a degradation in accuracy with increasing depth, making it ideal for distinguishing between toxic and non-toxic plant species, which may exhibit only minute visual differences. Noor et al. (2022) found that CNNs, such as ResNet-50, can be very effective when combined with support vector machines (SVMs) in predicting poisonous plant species. Combining features learned by CNNs with the classification power of SVMs produced better results, demonstrating how method combinations can lead to more robust outcomes in plant classification.

Furthermore, Chamidullin (2022) explored the use of fine-grained recognition techniques that utilized side information to enhance the classification process. When more information, such as the habitat or growth conditions of a plant, accompanies the image, the model can adjust its output accordingly. This shows that when toxic and non-toxic plants appear visually similar, deep learning models that incorporate contextual data may yield more accurate classifications.

Leaf patterns and venation are two features that can help identify whether a species is toxic or safe. Bhatt and Greenberg (2023) emphasized the importance of image-based toxicity classification, classifying plants based on their characteristics. Tasks like these are well-suited to the ResNet-50 model, as it can detect fine-grained visual details that allow it to differentiate between species based on subtle differences. This is particularly important for plant species in which slight variations in shape and texture indicate toxicity.

Color is another crucial characteristic for plant classification. Toxic plants often exhibit distinct color patterns on their leaves, flowers, or stems as indicators of toxicity. The ResNet-50 model, with its deep architecture, can classify plant species by capturing complex color variations at different levels of abstraction. Researchers have even attempted to identify and classify medicinal plants using deep learning models. Prasad (2024) discussed the use of deep learning models, emphasizing that color features contribute significantly to the accuracy of such systems. ResNet-50 is particularly well-suited to handling high-dimensional image data and performs exceptionally well when minute color differences must be identified.

Recent advancements in data augmentation and transfer learning have further improved models like ResNet-50 for plant classification tasks. By artificially expanding the training dataset through techniques such as rotation, scaling, and flipping, researchers can increase the model's ability to generalize to unseen plant species. This is especially beneficial in toxic plant classification, where large, well-labeled datasets are often unavailable. Alobeidli et al. demonstrated the effectiveness of transfer learning in classifying toxic and non-toxic plant species in the UAE by fine-tuning a pre-trained ResNet-50 model on a smaller, domain-specific dataset. Transfer learning reduces the need for extensive labeled data while maintaining high accuracy, making it a practical approach for real-world applications.

When applying the ResNet-50 model to real-world scenarios, it is likely that the model will encounter plant species it has not seen during training. Fine-tuning the model with domain-specific data allows it to retain the robustness of the pre-trained weights while adapting to the peculiarities of poisonous and non-poisonous plants in different regions. This is important for global biodiversity, as plant species can vary significantly in appearance depending on their geographic location. Jahan et al. (2023) demonstrated a similar approach for categorizing toxic frogs, revealing that achieving high precision requires regional fine-tuning. This approach can be applied to plant classification, where pre-trained models like ResNet-50 serve as a starting point for specific applications across different ecosystems.

Real-time applications, such as mobile-based plant identification systems, are gaining more attention. Using lightweight versions of deep learning models or model compression, it is possible to build systems that are both accurate and computationally efficient. This could have relevant applications in agriculture, forestry, and environmental monitoring, where rapid and reliable identification of toxic plants can prevent harm to humans and animals. The scalability of ResNet-50 for such tasks, and its proven ability to classify fine-grained details, makes it well-suited for practical use in these areas.

The application of the pre-trained ResNet-50 model enables it to inherit the depth and accuracy of advanced neural networks with the flexibility of transfer learning and hybrid models. Studies by Azadnia et al. (2024) and Noor et al. (2022) have demonstrated that the model can efficiently identify plant species based on visual features such as leaf texture, shape, and color.

3 Proposed Model

On the basis of literature survey, we proposed a model for classification of toxic and non-toxic plants using ResNet-50 model that is shown in figure 1.

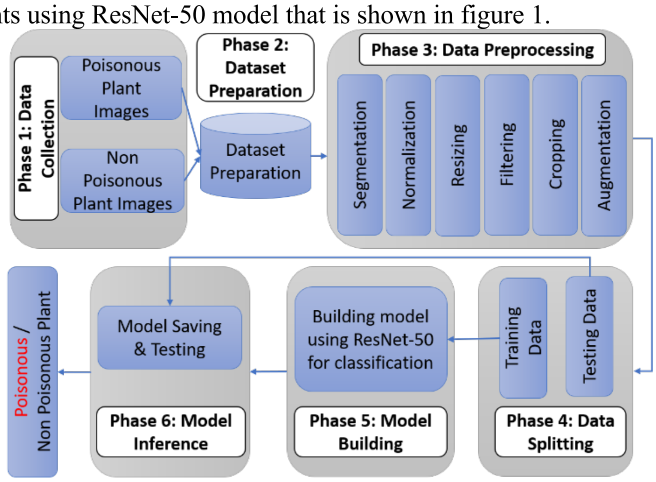

Figure 1: Proposed model

3.1 Phase 1: Data Collection

At this stage, images of toxic and non-toxic plants are downloaded from public datasets. The dataset contains an array of plant images to be classified accurately. Each image is labeled based on whether it contains a toxic plant or not. The dataset ensures sufficient diversity, quality, and quantity to support the classification task. This variation helps make the model more robust, enabling it to generate accurate results on unseen data as well.

3.2 Phase 2: Dataset Preparation

After gathering the data, we clean the dataset by organizing and structuring it properly for further processing. This involves carefully sorting and labeling the images of toxic versus non-toxic plants. A structured directory, with folders separating toxic and non-toxic plants, is developed. The dataset is then cleaned by removing duplicate entries, incomplete records, and any mislabeled images, ensuring that the dataset is accurate and ready for further processing.

3.3 Phase 3: Data Preprocessing

Once the dataset is ready, we begin the data preprocessing stage. In this stage, we employ several methods to prepare the data for training the model. Segmentation is used to isolate the plant from the background in the images, improving the focus on relevant features. This allows the model to better understand the plant's features while ignoring irrelevant background noise.

Next, we normalize the pixel intensities of the images to fall within a universal scale, typically between 0 and 1. This normalization ensures stable learning and results in faster convergence during training. The images are resized to 224x224 pixels to align with the input size required by ResNet-50, making all images uniform in size.

Additionally, filtering and cropping methods are applied to eliminate unwanted elements, ensuring the plant remains the focal point of the image. Another critical element of this step is data augmentation, where the dataset is artificially expanded by applying transformations such as rotations, flips, zooms, and shifts. These augmentations make the model more robust to variations in the input images, enhancing its generalization capability in real-world environments.

3.4 Phase 4: Data Splitting

After preprocessing the dataset, we divide it into training and testing data. This is done to assess the actual performance of the model. Typically, the dataset is split with 80% allocated for training and 20% for testing. The majority of the images are in the training set, allowing the model to learn from a wide variety of examples, while the testing set is reserved for evaluating the model's performance on unseen data.

During the split, we ensure that both the poisonous and nonpoisonous plant categories have equal representation in the training and testing phases to avoid any bias toward one category. This ensures the model's performance is as balanced as possible. The test set represents real-world data, simulating the behavior expected of the model when it encounters new images of plants after deployment.

3.5 Phase 5: Model Building

In this stage, we rely on the ResNet-50 pre-trained model. As a well-established deep convolutional neural network, ResNet-50 works effectively with transfer learning. We fine-tune the pre-trained ResNet-50 model by replacing its top layers with a custom classifier tailored to the task of toxic and non-toxic plant classification. This approach reduces training time and increases accuracy because ResNet-50 has already learned useful features from a large amount of image data during its initial training on ImageNet.

We freeze the early layers of ResNet-50 to preserve the learned low-level features, such as edges, shapes, and textures. We then train the upper layers of the network on our dataset, allowing the model to adapt specifically to distinguishing poisonous from non-poisonous plants. We also apply dropout layers and regularization techniques to prevent the model from overfitting.

The model uses a hybrid loss function, combining categorical cross-entropy with Adam optimization for efficient gradient descent. During the training process, we track key metrics such as accuracy, precision, recall, and F1 score to ensure that our model generalizes well.

3.6 Phase 6: Model Inference

Once the model is trained and tested, we save the final version for inference purposes. The saved model is used to predict whether a given plant image is toxic or non-toxic. During inference, the model takes in new, unseen images and classifies them based on the features learned during the training phase.

The inference process works by passing the input image through the model, which then produces a probability score for each class: toxic or non-toxic. After applying a threshold, the final decision on whether a plant is toxic or not is made. This phase is heavily optimized so that predictions can be made quickly and efficiently, even when working with high-resolution images.

The accuracy of future inferences depends largely on how well the model was trained and tested in previous phases. Ensuring robust preprocessing, training, and evaluation will result in a well-generalized model that performs reliably in real-world scenarios.

4 Pseudocode: Proposed Model

The pseudocode for the above proposed model is shown below.

```
BEGIN
 # Phase 1: Data Collection
 FUNCTION collect_data():
   data ← collect_images_from_sources()
   label_data(data, "poisonous", "non-poisonous")
   RETURN data

 # Phase 2: Dataset Preparation
 FUNCTION prepare_dataset(data):
   organized_data ← organize_data_into_folders(data)
   cleaned_data ← check_and_remove_duplicates(organized_data)
   resized_data ← resize_images(cleaned_data, 224x224)
   RETURN resized_data

 # Phase 3: Data Preprocessing
 FUNCTION preprocess_data(data):
   segmented_data ← segment_images(data)
   normalized_data ← normalize_pixel_values(segmented_data, 0, 1)
   augmented_data ← augment_data(normalized_data)
   RETURN augmented_data

 # Phase 4: Data Splitting
 FUNCTION split_data(data):
   training_data, testing_data ← split(data, ratio=0.8)
   RETURN training_data, testing_data

 # Phase 5: Model Building (ResNet-50 Pretrained)
 FUNCTION build_model():
   model ← load_pretrained_resnet50()
   freeze_initial_layers(model)
   add_custom_classification_layer(model, output_classes=2)
   compile_model(model,                loss_function="categorical_cross_entropy",
optimizer="Adam")
   RETURN model


 FUNCTION train_model(model, training_data):
   model ← train(model, training_data, epochs=NUM_EPOCHS)
   RETURN model

FUNCTION evaluate_model(model, testing_data):
   accuracy ← calculate_accuracy(model, testing_data)
   precision, recall, f1_score ← calculate_precision_recall_f1(model, testing_data)
   confusion_matrix ← generate_confusion_matrix(model, testing_data)
   IF performance_needs_improvement:
```

```
    fine_tune_model(model)
  RETURN accuracy, precision, recall, f1_score

# Phase 6: Model Inference
FUNCTION infer(model, new_image):
  prediction ← model.predict(new_image)
  IF prediction_score(poisonous) > threshold:
    RETURN "Poisonous"
  ELSE:
    RETURN "Non-Poisonous"


# Main Workflow
data ← collect_data()
prepared_data ← prepare_dataset(data)
processed_data ← preprocess_data(prepared_data)
training_data, testing_data ← split_data(processed_data)

model ← build_model()
trained_model ← train_model(model, training_data)

evaluation_metrics ← evaluate_model(trained_model, testing_data)

END
```

5 Result & Discussion

5.1 Dataset Description

The "Toxic Plant Classification" dataset is used for the implementation of the proposed model. It contains a diverse collection of plant images, specifically focused on the classification of toxic and non-toxic plants. The dataset includes 4,928 high-resolution images, divided into two main classes: (i) toxic and (ii) non-toxic plants, making it suitable for training the models. This dataset is well-structured, offering a balanced number of images across both classes, ensuring that models trained on it can generalize well without being biased toward a particular class.

5.2 Performance Metrics

For any classification problems, performance metrics play an important role in evaluating the effectiveness of a model. These metrics help assess how well the model is performing, not only in terms of overall accuracy but also in understanding the model's ability to generalize, handle imbalanced classes, and make reliable predictions across different classes.

Four parameters of confusion matrix; True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN); are used for the evaluation of model. Details of these parameters are given below:

1. TP: Both actual and model output classified are positive.
2. TN: Both actual and model output classified are negative.
3. FP: Actual classification is negative and model classified as positive.
4. FN: Actual classification is positive and model classified as negative.

The performance metrics based on above parameters are explained below:

1. Accuracy: It is the ratio of correctly predicted instances to the total number of instances.
2. Precision: It is the proportion of true positive predictions (correctly predicted positive instances) to the total positive predictions made by the model (true positives + false positives).
3. Recall (Sensitivity or True Positive Rate): It is the ratio of correctly predicted positive observations to all actual positive observations.
4. F1-Score: It is the harmonic mean of precision and recall. It provides a single metric that balances the two, especially in cases of imbalanced datasets.

5.3 Model result

The confusion matrix provides a detailed breakdown of the classification model's performance across four key outcomes: TP, TN, FP and FN.

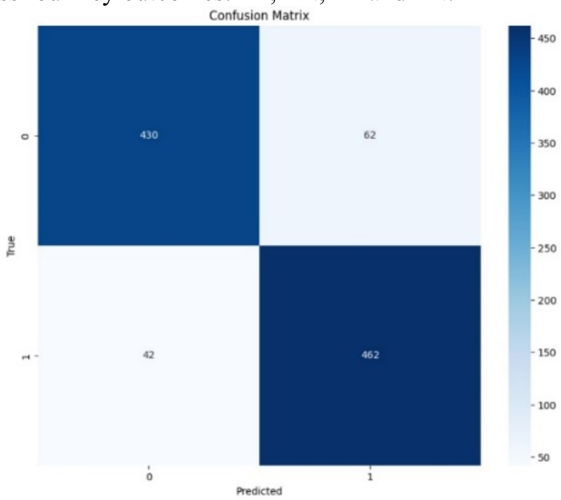

Figure 2: Confusion Matrix

Figure 2 showing confusion matrix that comes after the execution of proposed model on our dataset. Based on the above matrix the performance metrics are given below:

1. **Accuracy:** Total number of records used for testing are 996 images and number of records for TP and TN are 430 and 462 images respectively.

$$Accuracy = \frac{430 + 462}{430 + 62 + 42 + 462} = 0.896\ (or\ 89.6\%)$$

This indicates that the model correctly classifies approximately 89.6% of the instances.

2. **Precision:** Total number of TP images are 430 and FP images are 62. Precision of proposed model is;

$$Precision = \frac{430}{430 + 62} = 0.874\ (or\ 87.4\%)$$

This shows that 87.4% of the instances predicted as positive are actually positive.

3. **Recall:** Total number of TP images are 430 and FN images are 42. Recall of proposed model is;

$$Recall = \frac{430}{430 + 42} = 0.911\ (or\ 91.1\%)$$

This indicates that the model identifies 91.1% of the actual positive instances.

4. **F1-Score:** F1 score is calculated by the following formula.

$$F1 = 2 * \frac{0.874 * 0.911}{0.874 + 0.911} = 0.892\ (or\ 89.2\%)$$

This score reflects a good balance between precision and recall, suggesting the model performs well overall for the positive class.

5.4 Discussion

The performance metrics achieved by the proposed model highlight its suitability for classifying plants into toxic and non-toxic categories. In this study, the model achieved 89.6% accuracy, 87.4% precision, 91.1% recall, and an 89.2% F1-score. These results demonstrate the model's high reliability in correctly identifying plant species, minimizing both false positives and false negatives. Precision is especially critical when dealing with toxic plants, as misclassification could have serious consequences, including potential harm to humans or animals (Dwivedi et al., 2021). The high recall ensures that the model effectively identifies the majority of toxic plants, offering a robust solution for real-world applications such as mobile plant identification tools or agricultural safety systems (Azadnia et al., 2024).

6 Conclusion & Future work

In conclusion, the application of ResNet-50 for the classification of toxic and non-toxic plants offers a reliable, efficient, and scalable solution to a long-standing challenge in public safety and environmental monitoring. The model's ability to accurately differentiate between these plant categories, combined with its strong performance metrics, underscores its potential for integration into practical, real-world applications. Future work in this domain should focus on enhancing model accuracy through larger datasets, hybrid architectures, and the exploration of other advanced neural networks. Such developments will significantly benefit fields like agriculture, healthcare, and ecological conservation, providing accessible tools for the timely identification and management of plant species.

References

1. Azadnia, R., Noei-Khodabadi, F., Moloudzadeh, A., Jahanbakhshi, A. and Omid, M., 2024. Medicinal and poisonous plants classification from visual characteristics of leaves using computer vision and deep neural networks. Ecological Informatics, p.102683.
2. Zuhri, M.F., Maharani, S.K.R., Affandy, A., Nurhindarto, A., Syukur, A. and Soeleman, M.A., 2022. Classification of Toxic Plants on Leaf Patterns Using Gray Level Co-Occurrence Matrix (GLCM) with Neural Network Method. Journal of Development Research, 6(1), pp.1-5.
3. Hassoon, I.M. and Hantoosh, S.A., 2023. CFNN for Identifying Poisonous Plants. Baghdad Science Journal, 20(3 (Suppl.)), pp.1122-1122.
4. Mezzasalma, V., Ganopoulos, I., Galimberti, A., Cornara, L., Ferri, E. and Labra, M., 2017. Poisonous or non-poisonous plants? DNA-based tools and applications for accurate identification. International journal of legal medicine, 131, pp.1-19.
5. Dwivedi, H., Bhardwaj, M. and Kumar, G., 2021. A Comprehensive Study of Poisonous Plants of Family Apiaceae. Apiaceae: Ecology, Uses and Toxicity. Nova Publisher, New York, pp.1-14.
6. Wilk, I., Wnuk, E., Stachurska, A., Janicka, W., Tkaczyk, E., Kumanowicz, N. and Łuszczyński, J., 2024. Explorative behaviour in horses when presented with unfamiliar poisonous and non-poisonous plants. Applied Animal Behaviour Science, 277, p.106352.
7. Wendt, S., Lübbert, C., Begemann, K., Prasa, D. and Franke, H., 2022. Poisoning by plants. Deutsches Ärzteblatt International, 119(18), p.317.
8. Noor, T.H., Noor, A. and Elmezain, M., 2022. Poisonous Plants Species Prediction Using a Convolutional Neural Network and Support Vector Machine Hybrid Model. Electronics, 11(22), p.3690.
9. H, S.A., Joshi, S.V. and N, H.C. 2024. Review on Poisonous Plants Detection Using Machine Learning. International Journal of Advanced Research in Science, Communication and Technology. doi: 10.48175/ijarsct-15328.
10. Prasad, R.M.D. 2024. Identification and Classification of Medicinal Plants using Deep Learning. Indian Scientific Journal Of Research In Engineering And Management. doi: 10.55041/ijsrem29221.
11. Azadnia, R., Noei-Khodabadi, F., Moloudzadeh, A., Jahanbakhshi, A. and Omid, M. 2024. Medicinal and poisonous plants classification from visual characteristics of leaves using computer vision and deep neural networks. Ecological Informatics 82, p. 102683. doi: 10.1016/j.ecoinf.2024.102683.
12. Zuhri, M.F., Maharani, S.K.R., Affandy, A., Nurhindarto, A., Syukur, A. and Soeleman, M.A., 2022. Classification of Toxic Plants on Leaf Patterns Using Gray Level Co-Occurrence Matrix (GLCM) with Neural Network Method. Journal of Development Research, 6(1), pp.1-5.
13. Noor, T.H., Noor, A. and Elmezain, M., 2022. Poisonous Plants Species Prediction Using a Convolutional Neural Network and Support Vector Machine Hybrid Model. Electronics, 11(22), p.3690.
14. Chamidullin, R., 2022. Fine-grained Visual Recognition with Side Information.
15. Alobeidli, K., Shatnawi, M. and Almansoori, B., 2023, February. Image Classification for Toxic and Non-Toxic Plants in UAE. In 2023 Advances in Science and Engineering Technology International Conferences (ASET) (pp. 1-5). IEEE.
16. Jahan, M., Samia, I.J., Younus, S.B. and Hakim, M.A., 2023, July. A Deep Learning Based Poisonous Frog Detection System. In 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT) (pp. 1-6). IEEE.