SXC JOURNAL
Volume 2

Leveraging the performance of Deep Learning Models for Corn Leaf Disease Diagnosis using DenseNet201 and Xception

Rajan Karmacharya[1], Prashant Aryal[2] and Prashant Giri[2]*
[1]Department of Computer Science, St., Xavier's College, Maitighar, Kathmandu, Nepal*
[2]Department of Computer Science, St., Xavier's College, Maitighar Kathmandu, Nepal
[2]Department of Computer Science, St., Xavier's College, Maitighar, Kathmandu, Nepal
*rajankarmacharya@sxc.edu.np**

ABSTRACT

Plant diseases cause large output decreases and financial losses, making them a major barrier to global food security, especially in developing nations like Nepal, where corn is a staple crop. Early and accurate detection is critical to mitigating these impacts and improving crop management. Traditional diagnostic methods, reliant on manual inspection, are often time-consuming, subjective, and impractical for large-scale agricultural applications. This paper explores the automatic categorization of corn leaf diseases using deep learning-driven Convolutional Neural Networks (CNNs), specifically DenseNet 201 and Xception architectures. Convolutional layers in these models learn and extract distinctive features automatically from images, enabling accurate and efficient classification of corn disease types. A freely accessible dataset comprising images of both healthy and diseased corn leaves was utilized, with data augmentation strategies used to enhance model generalization and robustness. Experimental results demonstrate that DenseNet201 achieved a test accuracy of 98.69%, outperforming Xception, which attained 96.61%. These results demonstrate the highlights of CNN-based approaches for scalable, non-invasive, and accurate disease detection in corn crops. The proposed method offers a viable tool to support precision agriculture and contribute to enhancing global food security.

Keywords: Corn Disease, Deep Learning, Convolutional Neural Networks (CNNs), DenseNet201, Xception, Image Classification

Introduction

Increased commercialization is taking place in Nepal's agricultural sector, particularly in the production of vegetables and other important commodities. Just 281,132 hectares are used for vegetable growing, which yields an average of 15.09 metric tons per hectare and around 3,962,383 metric tons annually. In addition to vegetables, the nation's total agricultural output is significantly influenced by the commercial cultivation of staple crops such rice, maize, wheat, and barley. More reliance on agricultural supplies, such as artificial fertilizers, chemical pesticides, plant nutrients and improved seeds, has resulted from the growing demand for both vegetables and staple crops. But this increased input utilization has also made the crops more susceptible to several diseases and other problems in agriculture (Ministry of Agriculture and Livestock Development, 2020).

Figure 1: Growing corn fields (Republica, 2023)

The prosperity of farmers is deeply connected to the health and productivity of their crops, making effective agricultural management essential. In Nepal, where maize cultivation covers 28% of agricultural land and contributes significantly to food security, crop diseases cause annual yield losses of 35-40%. As farming practices become increasingly commercialized and scaled, traditional methods of disease detection—reliant on manual observation—are often inadequate, leading to delayed or inaccurate diagnoses. This is particularly critical in the case of corn, a staple crop valued for both its economic importance and nutritional benefits. Corn diseases like Common Rust, Northern Leaf Blight, and Gray Leaf Spot pose significant threats, reducing yields and driving up production costs, which directly impact farmers' livelihoods.

Moreover, the intensified use of fertilizers and pesticides associated with modern agricultural practices can increase crop susceptibility to these diseases. This study aims to answer the following research questions:

- How can deep learning techniques contribute to automating the early detection of corn leaf diseases, especially in countries like Nepal?
- Which model is more effective, DenseNet201 or Xception, in detecting and classifying corn leaf diseases based on test accuracy and performance metrics?

The findings aim to support the development of scalable, automated diagnostic tools that help farmers protect their crops and improve agricultural productivity.

Literature Review

Zhang et al. (2018) explored the use of deep learning methods for the automatic detection of maize leaf diseases, a critical need in agricultural information systems. Their study proposed improved versions of GoogLeNet and Cifar10 models, designed to enhance identification accuracy while minimizing network complexity. The upgraded GoogLeNet architecture attained an average identification accuracy of 98.9%, compared to 98.8% for the Cifar10 model. However, while the study showcases the strength of deep learning in agricultural applications, its reliance on standard CNN architectures limits the scope of advancements. This study addresses the same problem but with more recent models like DenseNet201 and Xception, which offer better performance in terms of feature reuse and computational efficiency. The use of pre-trained networks and transfer learning in the current research further enhances these advancements.

Krizhevsky et al. (2017) aimed to advance image classification by leveraging a powerful CNN on the ImageNet dataset. The proposed model achieved top-tier results, outperforming previous methods, with top-5 error rates of 17.0% on ImageNet and 15.3% in the ILSVRC-2012 competition. Overcoming overfitting challenges, the study utilized data augmentation, dropout, and a highly efficient GPU implementation for training. The paper emphasizes the importance of depth in achieving superior performance and discusses the potential for further improvement with faster GPUs and larger datasets.

Lv et al. (2020) addressed the challenges associated with maize leaf disease identification, particularly the difficulty of extracting lesion features under variable environmental conditions and inconsistent lighting. To improve recognition accuracy, the authors proposed a novel approach integrating a feature enhancement framework with an advanced neural network architecture, DMS-Robust AlexNet, built upon the standard AlexNet model. This architecture incorporated multi-scale and dilated convolutions to enhance feature extraction capabilities, while batch normalization was applied to reduce overfitting and improve model robustness. Additionally, AdaBound optimizer and PReLU activation function were utilized to enhance convergence and classification accuracy. Experimental validation demonstrated that the proposed feature enhancement algorithm significantly improved the model's capacity to recognize maize leaf diseases, even in natural, uncontrolled environments. These findings highlight the potential of advanced convolutional techniques for intelligent plant disease diagnosis, providing valuable insights for the broader integration of deep learning in precision agriculture.

Hu et al. (2020) explained the implications of transfer learning and data augmentation methods to enhance the detection of corn leaf diseases by means of deep learning models. Recognizing the difficulty of diagnosing corn diseases through manual observation—often resulting in misdiagnosis and reduced production efficiency—the authors proposed a CNN framework that integrates data augmentation with transfer learning to enhance diagnostic accuracy and model generalization. In their study, PlantVillage dataset images were utilized to classify into four different categories namely Corn Common Rust, Corn Northern Leaf Blight, Corn Gray Leaf Spot, and healthy leaves. The optimized model was developed by refining the GoogLeNet pre-trained network and fine-tuning parameters like the learning rate and optimizer. The authors evaluated the result of their optimized model with several other transfer learning-based CNN architectures, including ResNet18, VGG16, and VGG19. The results demonstrated that their projected model achieved an average accuracy of 97.6%, with each disease category achieving over 95% accuracy. Notably, their optimized model improved the accuracy by 5.9% compared to the original GoogLeNet model.

Wu et al. (2020) presented a study that focused on the challenging task of scene image classification, utilizing transfer learning with the Xception model, a CNN acknowledged for its noteworthy results in image classification. The research was conducted using the Intel Image Classification dataset, comprising 25,000 images across six natural scene categories. An assimilation of pre-trained and randomized weights was applied to train and test the models, with a comparison to the Inception-V3 model. Experimental results showed that transfer learning approach with Xception gained a high accuracy of 91.20% on scene classification tasks, outperforming the Inception-V3 model in robustness and generalization. The study highlighted the potential of Xception for scene classification tasks while suggesting future exploration into the effects of image resolution and novel transfer learning techniques for improved performance.

Zhong et al. (2018) researched the use of the DenseNet architecture to improve the detection of metastatic cancer using small image patches extracted out of digital pathology scans. The study employs DenseNet201 and its augmented version (DenseNet201(TTA)) to tackle challenges in cancer classification. Utilizing the modified PatchCamelyon (PCam) dataset, which consists of 220,025 samples, the research focuses on binary classification, distinguishing cancerous from non-cancerous samples. Results show that the DenseNet model wins over other architectures like ResNet34 and VGG19, with an AUC-ROC score of 0.971 and an accuracy of 98.9%. The paper emphasizes the use of data augmentation, PyTorch for model implementation, and the Adam optimizer to enhance training and performance.

Fuentes et al. (2017) conducted a study that focused on addressing challenges in agriculture by developing a way to detect and recognizing diseases and pests in tomato plants. They evaluated three deep learning meta-architectures including Faster R-CNN, R-FCN, SSD in combination with feature extractors like VGG net and ResNet. Using the Tomato Diseases and Pests Dataset, gathered from actual setups with diverse conditions, the study demonstrated the system's ability to identify nine various types of pests and diseases in tomatoes, including complex variations. The proposed approach, capable of real-time processing, provides a real-world and appropriate explanation for agricultural disease identification and localization, with potential applications beyond tomatoes, contributing to the advancement of agricultural research.

Identifying Gaps and Advancements

While prior research has established the effectiveness of deep learning for plant disease classification, several critical gaps remain. Existing studies, such as those by Zhang et al. (2018) and Hu et al. (2020) relied on older CNN architectures like GoogLeNet or AlexNet, which lack the feature reuse efficiency of DenseNet (Huang et al., 2016) or the computational optimization of Xception (Chollet, 2017). Additionally, most works, including Lv et al. (2020) used datasets with limited variability in lighting, background, or regional disease manifestations, reducing real-world applicability. Furthermore, few studies addressed practical deployment challenges, such as computational resource constraints in developing regions like Nepal, with notable exceptions like Fuentes et al. (2017) focusing only on tomatoes and Wu et al. (2020) examining general scene classification. The field has seen significant advancements with DenseNet's dense connectivity pattern enabling better feature reuse and Xception's depthwise separable convolutions offering computational efficiency, presenting new opportunities to overcome these limitations in plant disease classification.

Research Methodology and Materials

This study employs deep learning techniques to assess the efficiency of DenseNet201 and Xception models for classifying corn diseases.

The proposed system goes through several steps, including dataset preparation, data preprocessing which includes augmentation, rescaling and resizing, then model architecture design, and finally, evaluating the model's performance through a variety of evaluation metrics and ensuring the system's reliability.

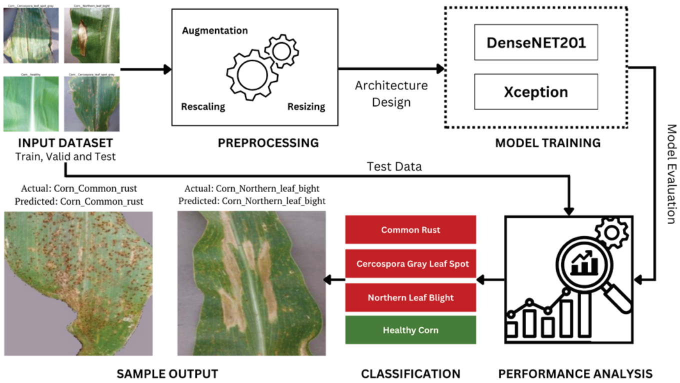

Figure 2: Proposed system stages for Corn Leaf disease diagnosis study.

Dataset Preparation

The dataset employed in this research work was obtained from Mendeley, a publicly accessible online repository and consists of corn leaf images, which were categorized into four classes: Common Rust, Northern Leaf Blight, Healthy Corn and Gray Leaf Spot. It is freely available for academic and research purposes, ensuring ethical compliance.
Dataset URL:
https://data.mendeley.com/datasets/tywbtsjrjv/1

While the dataset is balanced, potential limitations include regional bias (e.g., underrepresented geographic variations in disease manifestations) and variability in image quality (lighting, resolution), which may affect model generalizability.

The dataset was categorized into three different subsets: train, valid, and test. The training set consisted of 3,924 images, the validation set had 402 images, and the test set had 395 images. Each subset included images from all four categories, ensuring a balanced representation.

Figure 3: Sample images from multiple classes.

Table 1: Summary of the dataset

		Images			
	Class	**Train**	**Valid**	**Test**	**Total**
Diseased	Common Rust	944	99	101	1144
	Gray Leaf Spot	971	95	99	1165
	Northern Leaf blight	1010	101	101	1212
Healthy	Healthy Leaf	999	100	101	1200

Data Preprocessing and Augmentation

Data preprocessing and augmentation were employed to artificially expand and variety of the training dataset, addressing issues of class imbalance and reducing overfitting. The following augmentation techniques were applied:

- Resizing: A consistent dimension of 224 x 224 pixels was taken to resize the images, ensuring that all images in the dataset had the same input size, facilitating more effective training.
- Rescaling: Images were rescaled by dividing pixel values by 255 (i.e., 1.0/255) to normalize

them to a range of [0, 1], improving model performance.

After resizing and rescaling, the following augmentation techniques were applied to further increase the variability and robustness of the dataset:

- Rotation: Images were randomly rotated within a specified range of 15 degrees to simulate variations in leaf orientation.
- Flipping: Horizontal flipping was applied to simulate different perspectives.
- Shifting: Random width and height shifts within a range of 0.2 were used to mimic slight translations in the leaf positions.
- Zooming: Zooming was applied within a 0.2 range to simulate variations in proximity to the lens.
- Brightness Adjustment: Brightness was randomly adjusted within a range [0.95, 1.05] to account for different lighting conditions.

These augmentations enabled the model to learn more about the disease patterns, improving its ability to generalize with a robust representation.

Deep Learning and Convolutional Neural Network (CNN)

Deep learning, a branch of machine learning, has achieved significant success in domains like image classification through the use of deep neural network architectures. Among these, Convolutional Neural Networks (CNNs) have emerged as highly effective for visual tasks due to their ability to automatically learn hierarchical spatial features from raw images. CNNs consist of convolutional layers for feature extraction and pooling layers for dimensionality reduction, followed by fully connected layers for classification (MathWorks, 2024).

In the context of this study, CNNs serve as the backbone for corn leaf disease classification, enabling automated diagnosis by learning complex patterns from diseased and healthy leaf images. However, deep learning models typically require large amounts of data for effective generalization, which can be challenging when working with smaller or imbalanced datasets. To address this, we employ transfer learning and data augmentation, allowing pre-trained models like DenseNet201 and Xception to adapt effectively to the corn disease classification task with limited data (Simplilearn, 2023).

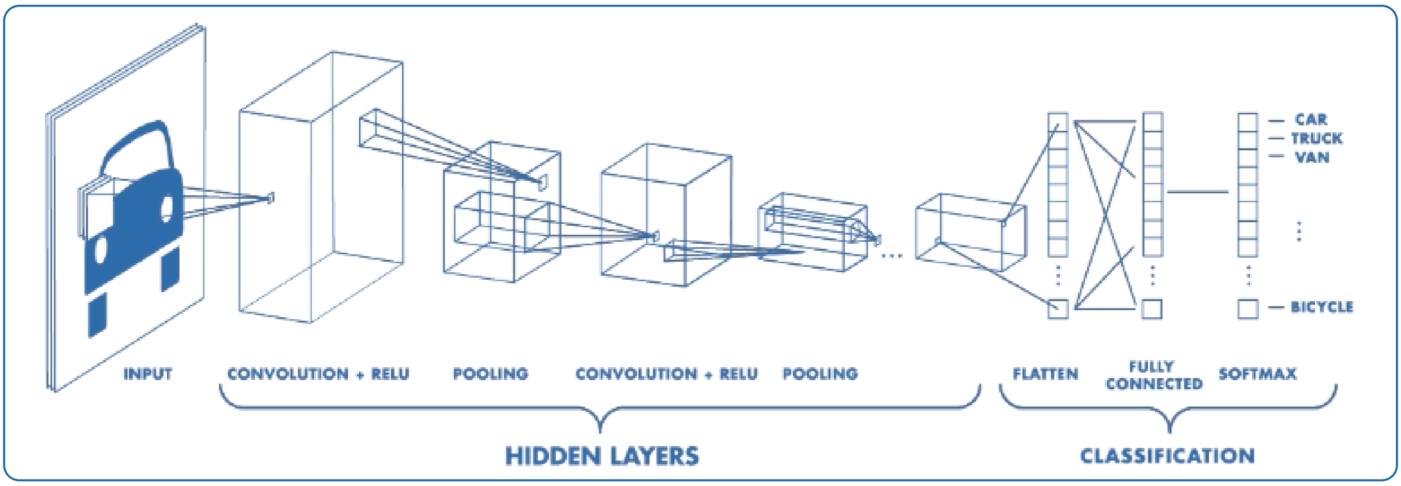

Figure 4: CNN Architecture. (Mishra, 2020)

The main layers of a CNN are as follows:

Convolutional Layer

This layer is a primary element of CNN, where most of the computation occurs. Filters (kernels) convolve over the input data to detect specific features in the image, typically represented as a 3D matrix (height, width, depth) for color images. The convolution operation scans the receptive fields for patterns (Ratan, 2020).

Pooling Layer

The feature map's spatial size is decreased through pooling. A filter slides over the feature map to select values from each receptive field. While average pooling calculates the average, max pooling chooses the maximum value (Seb, 2021).

Activation Layer
The ReLU activation function is widely deployed in CNNs. For negative inputs, it returns zero; for positive values, it returns the input itself. Its formula is f(x) = max(0, x) (Praharsha, 2024).

Fully Connected Layer
This layer connects all neurons, transforming input vectors via a weight matrix and applying a non-linear activation function. It ensures every input affects every output (Unzueta, 2022).

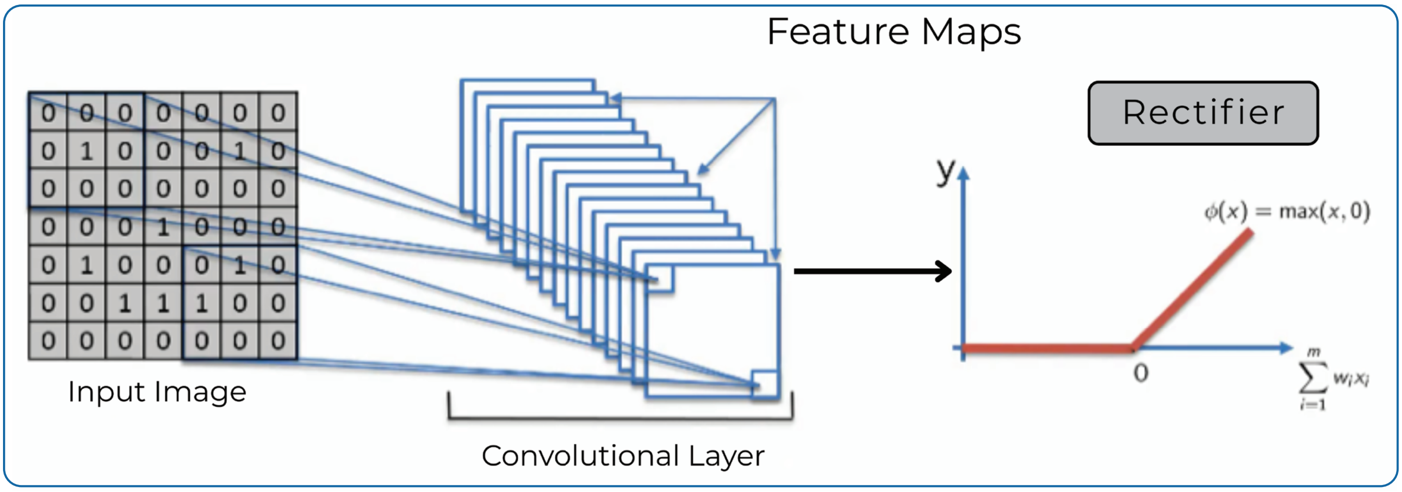

Figure 5: Activation Layer (SuperDataScience, 2018)

Normalization Layer
Normalization techniques like batch and layer normalization standardize inputs or outputs to stabilize training and speed up convergence (Bala, Build Better Deep Learning Models with Batch and Layer Normalization, 2023).

Dropout Layer
Dropouts randomly deactivate neurons during training to prevent overfitting. It encourages the network to generalize better by avoiding reliance on specific neurons (Brownlee, 2022).

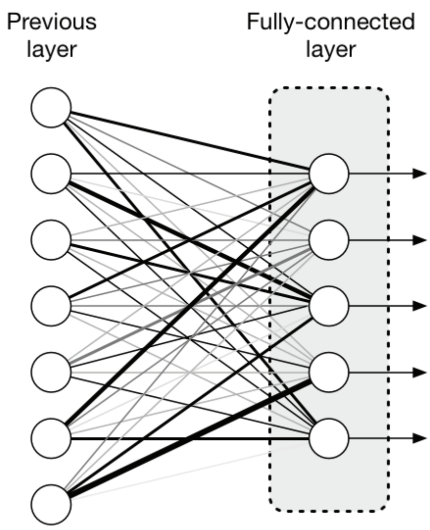

Figure 6: Fully Connected Layer. (Teco KIDS, 2019)

Dense Layer
The dense layer combines the features extracted from previous layers to make predictions. It typically follows the convolutional and pooling layers, using flattened data to output predictions (Verma, 2024).

Softmax Activation Function
Softmax converts network output into a distributed probability, which is particularly useful for multiclass classification tasks. Each value in the resulting vector indicates the likelihood that the input falls into a certain class, with the probabilities summing up to 1. This makes Softmax ideal for determining the most likely class for a given input, providing a clear decision boundary for classification (Bala, Softmax Activation Function, 2023).

In terms of mathematics, it is represented as:

$$s\left(x_i\right) = \frac{e^{x_i}}{\sum_{j=1}^{n} e^{x_j}}$$

Figure 7: Softmax Activation Function Formula. (Apple Developer, 2023)

Transfer Learning

Transfer learning is essential to Deep Learning because it makes it possible to apply knowledge from previously developed models to new problems, especially those with sparse data. It transforms the field of data science and improves its ability to handle challenging challenges by utilizing information generated by comparable jobs. With this approach, learned expertise is effectively transferred from one domain to another by using the insights acquired from an already trained model that enhances the predictions for a new task. In situations when getting labeled data has issues related to cost or accessibility, this approach saves time and computational resources as well as aids in the creation of machine learning systems that are more reliable and accurate (Sharma, 2023).

Model Architecture

For this study, DenseNet201 and Xception were selected for classifying corn diseases due to their proven success with small datasets. DenseNet201's densely connected layers enhance feature reuse and improve generalization, while Xception's depthwise separable convolutions reduce computational complexity, making it efficient for fine-grained tasks like corn leaf disease detection. Both models have shown success in similar plant disease classification tasks, making them well-suited for this study. Although models like EfficientNet, ResNet, and InceptionV3 perform well in large-scale tasks, EfficientNet's tendency to overfit small datasets, along with its higher computational demands and the complexity of ResNet and InceptionV3, made DenseNet201 and Xception a more suitable choice for this study and real-time agricultural applications with limited resources.

Here, ImageNet pretrained weights were used to initialize these models, and to improve them for particular purpose of identifying corn diseases, transfer learning was used. The application of pretrained weights allowed the models to leverage features learned from a wide and holistic dataset, reducing the need for extensive training from scratch and improving generalization, especially for the smaller, domain-specific dataset. The following common parameters were used across both architectures:

- Input Shape: The input to both models was resized to 224 × 224 ×3 (224 pixels in width and height, and 3 color channels for RGB).
- Batch Size: The batch size used during training was 32, meaning that 32 images were processed simultaneously before updating the model's weights.
- Training Epochs: Both models were trained for 12 epochs, meaning the models iterated over the entire training dataset 12 times.
- Learning Rate: A learning rate of 0.0001 was used for fine-tuning the models, optimizing the training process.

DenseNet201

DenseNet (Densely Connected Convolutional Networks) is a deep learning architecture that incorporates dense connectivity, allowing every layer to acquire input from all previous layers. During training, this dense connection structure enhances gradient flow and encourages feature reuse, hence mitigating the disappearing gradient issue in very deep networks. DenseNet uses fewer parameters than typical CNNs by recycling features across layers, while maintaining or enhancing performance. The key principle underlying DenseNet is that each layer contributes to the next layer, resulting in a more densely coupled model in which every output feature map is sent to all subsequent layers, allowing for fast learning at several levels of abstraction (Huang et al., 2016).

DenseNet201 is a particular version of DenseNet architecture, consisting of 201 layers arranged in dense blocks with multiple convolutional layers. The model's main strength is its use of dense connections, which enable it to preserve detailed feature representations throughout the network. This improves generalization, making DenseNet201 particularly effective for jobs like image classification and plant disease detection, where extracting complex, discriminative features from data is essential (MathsWorks, 2024).

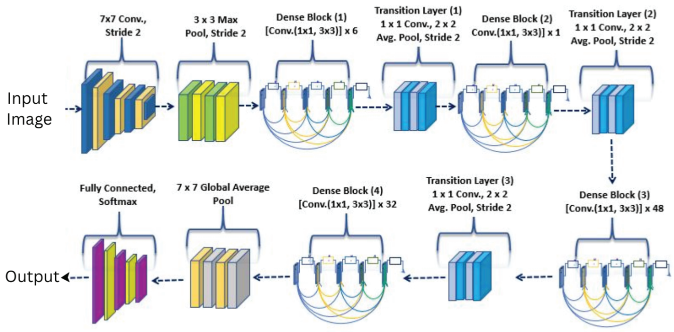

Figure 8: DenseNet201 Architecture (Akhtar, et al., 2022)

In this study, DenseNet201 was fine-tuned using transfer learning, where the lower layers were frozen with the last 20 layers were unfrozen to adapt the model to the corn disease dataset. Custom top layers were added, including Flatten, Batch Normalization, Dense layer (64 neurons), with L2 regularization applied to avoid overfitting, Dropout (0.6), and a final output layer with 4 neurons and Softmax activation.

Xception Architecture

Xception (Extreme Inception) is a CNN architecture that extends the Inception model by using depth-wise separable convolutions throughout the network. This technique splits the convolution process into two stages: pointwise convolution, which applies a 1x1 convolution across channels to integrate information, and depth-wise convolution, which processes each input channel individually. Xception maintains the model's capacity to identify intricate patterns by keeping these operations apart while significantly reducing both the cost of computation as well as number of parameters. As a result, Xception offers a more efficient and effective design, often resulting in enhanced performance while classifying images (Chollet, 2017).

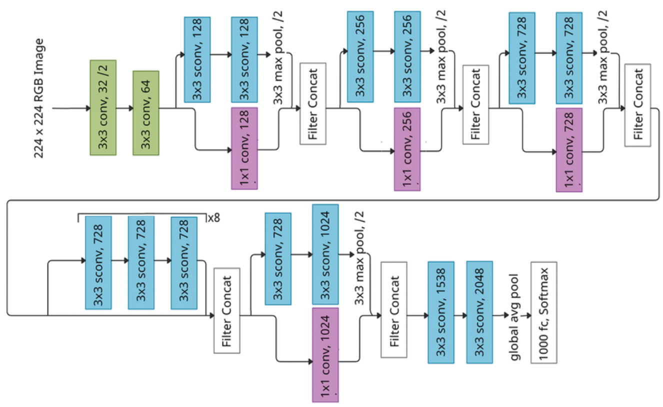

Figure 9: Xception Architecture. (Srinivasan, et al., 2021)

Similar to DenseNet201, Xception was fine-tuned using pre-trained weights from ImageNet, where the lower layers were frozen, keeping the higher layers unfrozen to adapt the model to the corn disease dataset. The model was modified by adding custom layers, including Flatten, Batch Normalization, Dense (64 neurons) with L2 regularization to prevent overfitting, Dropout (0.6), and a final Softmax output layer with 4 neurons for multiple-class classification.

Training Process

Adam optimizer, at a learning rate of 0.0001, was used to train the models. For multi-class classification problems with integer labels, the sparse categorical cross entropy loss function was employed, which works well. The models had a batch size of 32 and were trained for 12 epochs.

During training, data augmentation was applied to the training set in real-time through ImageDataGenerator to improve the variety of the dataset and avoid overfitting. A learning rate scheduler was employed to lessen the learning rate dynamically if the validation loss plateaued. This helped optimize the model's exactness over the course of training.

Evaluation Metrics

To guarantee a thorough analysis, the models' performance was evaluated using a variety of evaluation indicators. Accuracy was used to measure the proportion of correctly classified images, providing an overall performance indicator. Precision, F1-score and recall were calculated in order to assess the model's capacity to accurately identify and categorize each disease class. The confusion matrix used gave the distribution of true positives, false positives, true negatives, and false negatives, offering insight into misclassification patterns.

Additionally, ROC curves, for each class, were plotted and the AUC scores were calculated to evaluate the model's discriminatory power. In order to track the convergence of a model and identify overfitting during training, training and validation loss curves were studied. A classification report was generated to offer a thorough breakdown of each class's metrics. These metrics together provided a robust framework for evaluating the effectiveness of the DenseNet201 and Xception architectures.

Results and Analysis

The efficiency of DenseNet201 and Xception models was assessed with multiple metrics, which included accuracy, precision, recall, F1-score, and ROC-AUC, to analyze their usefulness in classifying corn diseases. The analysis also involved a comparison of training and validation curves and confusion matrices to evaluate the robustness and dependability of the models.

Model Performance

DenseNet201 classified corn diseases with 98.69% accuracy, demonstrating excellent performance. Its precision, recall, and F1-score were consistently above 98%, reflecting high reliability in identifying disease classes without significant misclassification. Xception achieved a slightly lower test accuracy of 96.61% but showed strong computational efficiency, making it suitable for applications requiring faster inference.

Table 2: Model Performance metrics on the test dataset.

Metric	DenseNet201	Xception
Accuracy	0.9869	0.9671
Precision	0.9878	0.9671
Recall	0.9873	0.9671
F1-Score	0.9873	0.9670

Confusion Matrix Analysis

Confusion matrices for both models highlight their ability to accurately classify the four disease classes with minimal misclassification. DenseNet201 misclassified only 5 samples out of 395, while Xception showed slightly higher misclassification, particularly in the Northern Leaf Blight classes and Gray Leaf Spot, but its overall performance remained strong across all disease categories.

The misclassifications could be due to visual similarities between diseases, especially in leaf patterns and discoloration, as well as dataset inconsistencies like image quality and lighting. Additionally, Xception's reliance on depthwise separable convolutions may limit its ability to capture fine details compared to DenseNet201, which uses dense connections to improve feature reuse and detail retention. This design helps DenseNet201 preserve richer information across layers, aiding in more accurate classification of subtle leaf disease differences. A more robust dataset and further fine-tuning could help reduce misclassifications in future iterations.

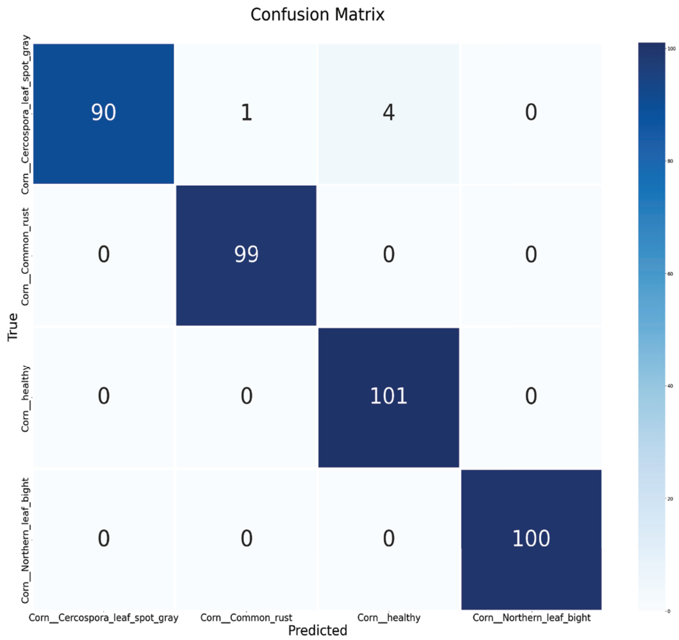

Figure 10: Confusion matrix for DenseNet201.

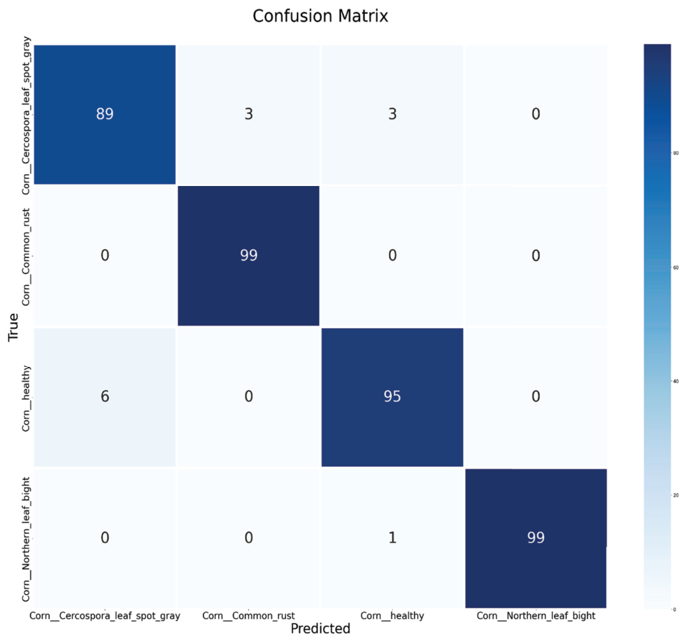

Figure 11: Confusion matrix for Xception.

Training and Validation Analysis

To see the behavior of models, training and validation curves were analyzed for precision and loss. DenseNet201 exhibited smooth convergence with minimal gaps between training and validation loss, indicating effective generalization to unseen data. There was a consistent improvement in training accuracy while the validation accuracy stabilized near the 98.78% mark by the final epoch, suggesting that the model generalized well.

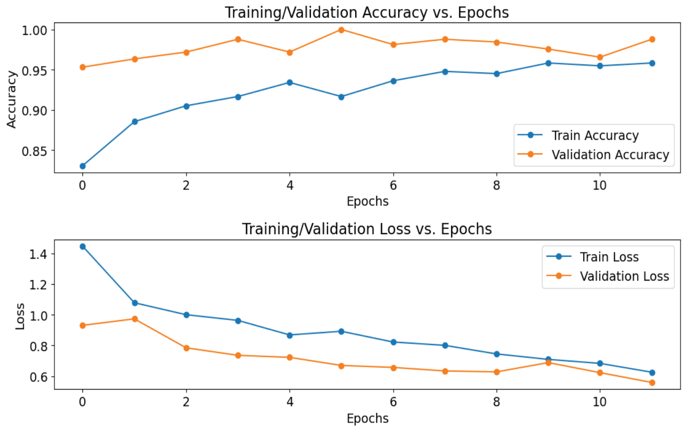

Figure 12: Training and validation accuracy and loss curves for DenseNet201.

In contrast, Xception exhibited a slightly higher validation loss during training, indicating some potential for improvement in generalization. The higher validation loss could be due to the model's complexity and the challenge of fine-tuning it for this specific dataset. Despite this, Xception still attained a notable validation accuracy of 97.56%, with training accuracy approaching 91.67% by the last epoch. This behavior suggests that while the model has room for improvement, it remains reliable and can be used for real-time applications that prioritize computational efficiency.

Figure 13: Training and validation accuracy and loss curves for Xception.

ROC Curve and AUC Analysis

The ROC curves for both models demonstrated excellent discrimination capabilities for all four classes. DenseNet201 achieved an AUC close to 1.0 for all classes. Similarly, Xception also achieved flawless AUC scores of 1.0 in every class of corn. These AUC values indicate that both models are highly reliable when distinguishing between disease classes, with DenseNet201 showing slightly better alignment with the top-left corner of the curve, signaling better prediction confidence. This makes DenseNet201 potentially more reliable in practical applications requiring high-confidence decisions.

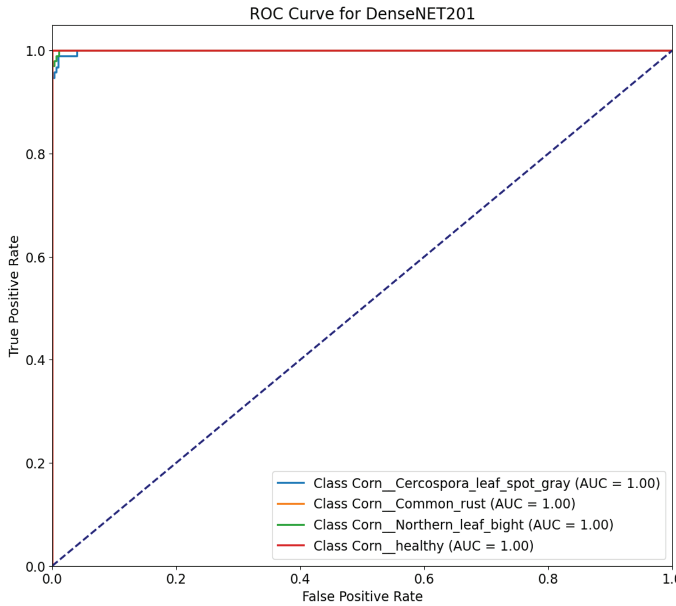

Figure 14: ROC curve and AUC score for DenseNet201.

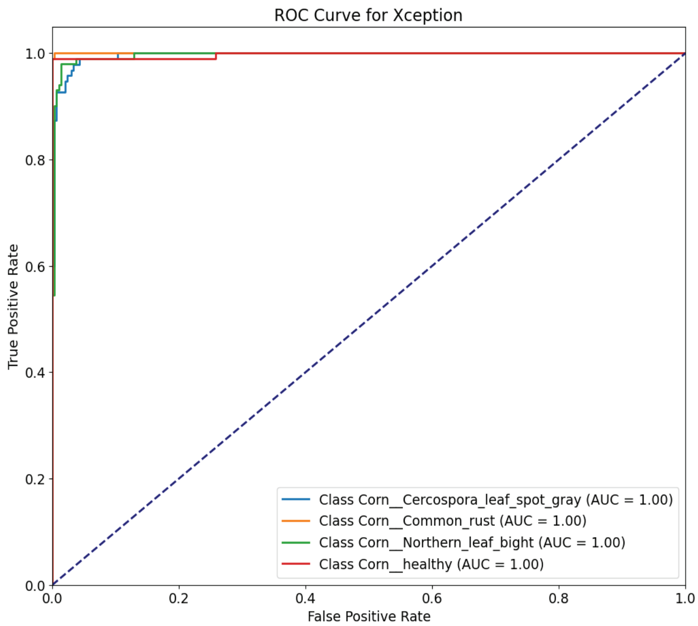

Figure 15: ROC curve and AUC score for Xception.

Model Predictions

To demonstrate the practical utility of the DenseNet201 model, several images collected from the test set were classified, and a comparison of the predicted labels against true labels was conducted. The model was chosen due to its higher accuracy, ensuring more reliable disease classification. The examples below highlight the ability of the model to correctly classify various corn diseases.

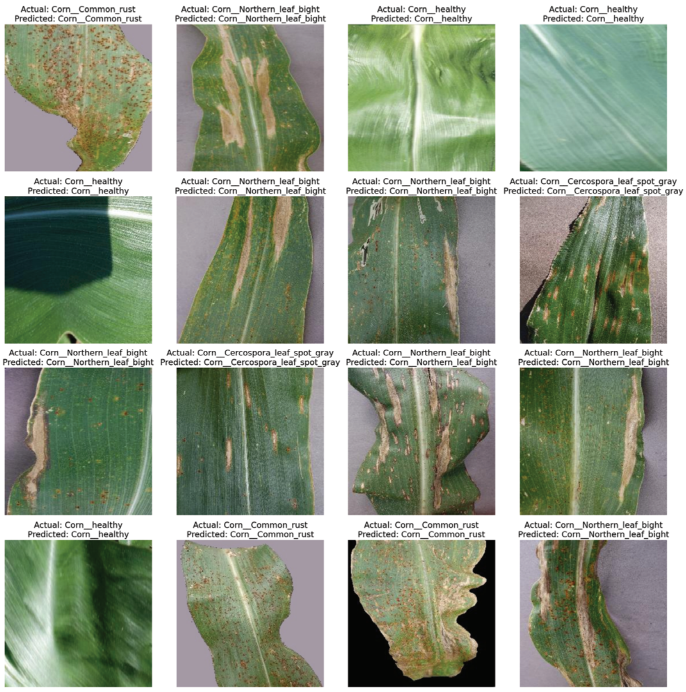

Figure 16: Sample classifications by DenseNet201.

Conclusion

This study demonstrates the efficacy of DenseNet201 and Xception architectures in the automated classification of corn leaf diseases, highlighting their potential for practical deployment in precision agriculture. DenseNet201 attained an imposing test accuracy of 98.69%, along with impressive recall, precision, and F1-scores, making it particularly effective for detecting complex disease patterns with high reliability. While Xception attained a slightly lower accuracy of 96.61%, it offers notable advantages in computational efficiency, enabling faster training and inference. The results confirm that convolutional neural networks (CNNs) provide a robust framework for scalable, non-invasive plant disease diagnosis. Furthermore, these findings underscore the importance of balancing model accuracy with computational efficiency when designing AI-driven solutions for crop health monitoring. The implementation of such models has the possibility to enhance early disease detection, optimize crop management practices, and eventually support greater food security and agricultural productivity. In practical terms, this approach could significantly reduce labor costs by automating disease identification and intervention, allowing farmers to detect and address issues earlier. This, in turn, would help minimize unnecessary pesticide use, leading to both cost savings and a reduction in the environmental impact caused by pesticide overuse.

Limitations and Recommendations

The performance of the model is limited by the small dataset, which affects its ability to generalize diverse real-world conditions and plant disease variations. Expanding the dataset with more varied images, including images from different geographical regions, would enhance the capacity of the model to differentiate between plant diseases and other objects. While DenseNet201 is highly accurate, it requires substantial computational resources, making it less practical to use in real-time settings with limited resources. To address this, optimizing the model through techniques like model pruning, or integrating DenseNet201 with a more efficient model like Xception, could help strike a better balance between accuracy and speed, making the solution more feasible for deployment in the field. Furthermore, developing deployment strategies that ensure affordability for farmers, including low-cost hardware solutions or cloud-based processing, would be beneficial.

Moreover, integrating IoT sensors with mobile apps for real-time disease monitoring could provide valuable data on disease progression, improving management. Additionally, enhancing the model's robustness to environmental factors, such as lighting and weather variations, would improve its real-world performance. Developing optimized mobile applications with explainable AI and transparent data management technologies, such as blockchain, could enhance reliability and accessibility for farmers, particularly in rural areas with limited resources, especially in developing regions like Nepal. To facilitate widespread adoption, pilot programs and field testing should be conducted to validate the model's real-world effectiveness and economic feasibility.

References

Akhtar, M. J., Mahum, R., Butt, F. S., Amin, R., El-Sherbeeny, A. M., Lee, S. M., & Shaikh, S. (2022, 10 22). A Robust Framework for Object Detection in a Traffic Surveillance System. *Electronics, 11*. doi:10.3390/electronics11213425

Apple Developer. (2023). *Softmax*. Retrieved February 05, 2024, from Apple Developer: https://developer.apple.com/documentation/accelerate/bnnsactivationfunction/2915301-softmax

Bala, P. C. (2023, June 23). *Build Better Deep Learning Models with Batch and Layer Normalization*. Retrieved January 12, 2025, from Pinecone: https://www.pinecone.io/learn/batch-layer-normalization/

Bala, P. C. (2023, 06 30). *Softmax Activation Function*. Retrieved January 13, 2025, from Pinecone: https://www.pinecone.io/learn/softmax-activation/?utm_content=260133499&utm_medium=social&utm_source=linkedin&hss_channel=lcp-20299330

Brownlee, J. (2022, August 06). *Dropout Regularization in Deep Learning Models with Keras*. Retrieved February 28, 2024, from Machine Learning Mastery: https://machinelearningmastery.com/dropout-regularization-deep-learning-models-keras/

Chollet, F. (2017, July 26). Xception: Deep Learning with Depthwise Separable Convolutions. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1800-1807. doi:10.1109/CVPR.2017.195

Fuentes, A., Yoon, S., Kim, S., & Park, D. (2017, September 4). A Robust Deep-Learning-Based Detector for Real-Time Tomato Plant Diseases and Pests Recognition. *Sensors in Agriculture, 17*(9). doi:https://doi.org/10.3390/s17092022

Hu, R., Zhang, S., Wang, P., Xu, G., Wang, D., & Qian, Y. (2020). The identification of corn leaf diseases based on transfer learning and data augmentation. *CSSE '20: Proceedings of the 3rd International Conference on Computer Science and Software Engineering*, (pp. 58-65). doi:https://doi.org/10.1145/3403746.3403905

Huang, G., Liu, Z., Weinberger, K. Q., & Maaten, L. v. (2016, August 25). Densely Connected Convolutional Networks. *CoRR, abs/1608.06993*. doi:https://doi.org/10.48550/arXiv.1608.06993

Krizhevsky, A., Sutskever, I., & Hinton, G. (2017). ImageNet Classification with Deep Convolutional Neural Networks. *Communications of the ACM, 60*(6), 84-90. doi:https://doi.org/10.1145/3065386

Lv, M., Zhou, G., He, M., Chen, A., Zhang, W., & Hu, Y. (2020). Maize Leaf Disease Identification Based on Feature Enhancement and DMS-Robust Alexnet. *IEEE Access, 8*, 57952-57966. doi:10.1109/ACCESS.2020.2982443

MathsWorks. (2024). *DenseNet-201 convolutional neural network*. Retrieved February 3, 2025, from MathWorks Deep Learning: https://www.mathworks.com/help/deeplearning/ref/densenet201.html

MathWorks. (2024). *What Is a Convolutional Neural Network?* Retrieved February 28, 2025, from Mathworks: https://www.mathworks.com/discovery/convolutional-neural-network.html

Ministry of Agriculture and Livestock Development. (2020). *Statistical Information on Nepalese Agriculture: Brief Information about Nepalese Agriculture*. Kathmandu, Nepal: Government of Nepal. Retrieved December 14, 2024, from https://moald.gov.np/wp-content/uploads/2022/04/STATISTICAL-INFORMATION-ON-NEPALESE-AGRICULTURE-2076-77-2019-20.pdf

Mishra, M. (2020, August 26). *Towards Data Science*. Retrieved February 28, 2024, from Convolutional Neural Networks, Explained: https://towardsdatascience.com/convolutional-neural-networks-explained-9cc5188c4939/

Pandian, A., & Geetharamani, G. (2019, April 18). *Data for: Identification of Plant Leaf Diseases Using a 9-layer Deep Convolutional Neural Network*. Retrieved November 16, 2024, from Mendeley Data: https://data.mendeley.com/datasets/tywbtsjrjv/1

Praharsha, V. (2024, March 17). *ReLU (Rectified Linear Unit) Activation Function*. Retrieved February 01, 2025, from OpenGenus: https://iq.opengenus.org/relu-activation/

Ratan, P. (2020, 10). *What is the Convolutional Neural Network Architecture?* Retrieved 02 01, 2024, from Analytics Vidhya: https://www.analyticsvidhya.com/blog/2020/10/what-is-the-convolutional-neural-network-architecture/

Republica. (2023, May 14). *Drought hits maize crop in Udayapur*. Retrieved December 11, 2024, from https://myrepublica.nagariknetwork.com/news/drought-hits-maize-crop-in-udayapur

Seb. (2021, December 1). *What is Pooling in a Convolutional Neural Network (CNN): Pooling Layers Explained*. Retrieved February 01, 2025, from Programmathically: https://programmathically.com/what-is-pooling-in-a-convolutional-neural-network-cnn-pooling-layers-explained/

Sharma, P. (2023, December 07). *Understanding Transfer Learning for Deep Learning*. Retrieved March 04, 2024, from https://www.analyticsvidhya.com/blog/2021/10/understanding-transfer-learning-for-deep-learning/

Simplilearn. (2023, March 26). *What is Deep Learning?* Retrieved December 09, 2024, from Simplilearn: https://www.simplilearn.com/tutorials/deep-learning-tutorial

Srinivasan, K., Alaboudi, A., Garg, L., Datta, D., Agarwal, R., Jhanjhi, N. Z., & Thomas, A. G. (2021). Performance Comparison of Deep CNN Models for Detecting Driver's Distraction. *Computers, Materials & Continua*. doi:https://doi.org/10.32604/cmc.2021.016736

SuperDataScience. (2018, August 17). *Convolutional Neural Networks (CNN): Step 1(b) - ReLU Layer*. Retrieved February 02, 2024, from https://www.superdatascience.com/blogs/convolutional-neural-networks-cnn-step-1b-relu-layer

Teco KIDS. (2019, 07 31). *Fully-Connected Layer with dynamic input shape*. Retrieved January 21, 2025, from Medium: https://medium.com/@tecokids.monastir/fully-connected-layer-with-dynamic-input-shape-70c869ae71af

Unzueta, D. (2022, 10 18). *Fully Connected Layer vs. Convolutional Layer: Explained*. Retrieved 02 01, 2024, from https://builtin.com/machine-learning/fully-connected-layer

Verma, Y. (2024, February 28). *Dense Layers in Neural Networks*. Retrieved January 15, 2025, from Analytics India Mag: https://analyticsindiamag.com/a-complete-understanding-of-dense-layers-in-neural-networks/

Wu, X., Liu, R., Yang, H., & Chen, Z. (2020). An Xception Based Convolutional Neural Network for Scene Image Classification with Transfer Learning. *2020 2nd International Conference on Information Technology and Computer Application (ITCA)* (pp. 262-267). IEEE. doi:10.1109/ITCA52113.2020.00063

Zhang, X., Qiao, Y., Meng, F., Fan, C., & Zhang, M. (2018). Identification of Maize Leaf Diseases Using Improved Deep Convolutional Neural Networks. *IEEE Access, 6*, 30370-30377. doi:10.1109/ACCESS.2018.2844405

Zhong, Z., Zheng, M., Mai, H., Zhao, J., & Liu, X. (2020, November 28). Cancer image classification based on DenseNet model. *Journal of Physics: Conference Series, 1651*. doi:10.1088/1742-6596/1651/1/012143