EdgeCutMix Augmentation: Enhancing the Leaf Disease Classification for the Minority Class

Derisma [1,2], Nur Rokhman [1,*], and Dyah Aruming Tyas [1]

[1] Department of Computer Science, Universitas Gadjah Mada, Yogyakarta, Indonesia
[2] Faculty of Information Technology, Universitas Andalas, Padang, Indonesia
Email: derisma@fti.unand.ac.id (D.); nurrokhman@ugm.ac.id (N.R.); dyah.aruming.t@ugm.ac.id (D.A.T.)
*Corresponding author

***Abstract*—Leaf disease classification faces significant challenges due to dataset imbalances, particularly within minority classes, leading to decreased model accuracy. This study addresses this problem by introducing EdgeCutMix, a novel image augmentation technique designed to enhance the representation of minority classes. EdgeCutMix integrates edge detection, using the Canny edge detection algorithm, and selective mixing strategies to generate realistic and informative augmented images. The Plant Pathology 2020 dataset, consisting of 3,642 apple leaf images, was used for evaluation. The experimental setup involved oversampling and comparison against existing techniques like MixUp, CutOut, CutMix, and Mosaic, and training on four CNN architectures: MobileNetV2, EfficientNetB7, ResNet50, and DenseNet201. Results showed that EdgeCutMix significantly improved classification accuracy for minority classes, achieving up to 98% accuracy with the EfficientNetB7 model. These findings suggest that EdgeCutMix provides a promising solution for improving model performance in imbalanced datasets, with potential applications in advancing deep learning in agricultural pathology.**

***Keywords*—leaf disease classification, image augmentation, minority class, imbalanced dataset, EdgeCutMix**

I. Introduction

The detection and classification of leaf diseases are critical in agriculture, directly impacting crop yields and resource management [1–3]. Deep learning, particularly Convolutional Neural Networks (CNNs), has emerged as a formidable tool in this domain due to its ability to process and analyze complex image data [4–6]. However, a significant challenge in this field is the imbalanced nature of datasets, where some diseases are substantially underrepresented [7, 8]. This imbalance often leads to biased models that perform well on common diseases but poorly on rare ones, which can be devastating given their potential impact on crops [7]. This phenomenon has been thoroughly discussed in the literature, highlighting the need for strategies to manage dataset imbalances [7–9].

For instance, studies have shown that traditional machine learning models trained on imbalanced datasets tend to have inflated overall accuracy metrics, which mask poor predictive performance on minority classes. This situation underscores the need for techniques that specifically address data imbalance to ensure robust, accurate disease classification across all classes, as discussed in the literature on learning from imbalanced data [7, 10].

In response to this challenge, our study introduces EdgeCutMix, an innovative data augmentation technique that enriches the training dataset by intelligently blending images based on edge detection and selective mixing. Unlike traditional augmentation methods like rotation and flipping, which do not necessarily address class imbalances, EdgeCutMix specifically enhances the representation of underrepresented classes, thus mitigating the model's bias towards more frequent conditions. This approach draws inspiration from proven techniques in the literature, such as CutMix and MixUp, which have been noted for their efficacy in model performance enhancement across varied scenarios [11, 12].

EdgeCutMix distinguishes itself from existing methods through its unique integration of edge detection algorithms, which ensure that critical features of the diseases are preserved during the augmentation process. This is crucial for maintaining the diagnostic integrity of the images. The technique's novelty also lies in its ability to blend features from different classes strategically, enhancing the dataset's diversity without compromising the quality of the training data. Theoretically, EdgeCutMix is grounded in the principles of both edge enhancement and region-based mixing, drawing from the strengths of methods like CutMix and MixUp but optimizing them to handle the nuances of plant disease imagery [11, 12].

By addressing these specific challenges and filling a notable gap in current augmentation methodologies, EdgeCutMix contributes to advancing the field of plant pathology through deep learning. This paper will detail the algorithm's implementation, supported by a mathematical framework that illustrates how EdgeCutMix modifies the training process to achieve a more balanced and representative dataset, ultimately leading to improved model performance and reliability in real-world agricultural settings.

Manuscript received April 1, 2024; revised June 6, 2024; accepted August 1, 2024; published November 27, 2024.

The remainder of this paper is organized as follows: Section II reviews related work highlighting the gap that EdgeCutMix aims to fill. Section III describes the EdgeCutMix algorithm, including its theoretical underpinnings and implementation steps. Section IV presents the experimental setup, datasets, and metrics used for evaluation. Section V offers a comparative analysis of EdgeCutMix against benchmark techniques. Finally, Section 6 concludes the paper with a summary of findings and potential directions for future research.

II. Literature Review

One way to overcome sample imbalance in leaf disease classification using CNN is through data augmentation, which involves image manipulation to create additional variation in training data, enriching the minority class [8]. Augmentation techniques such as translation, rotation, flipping, brightness adjustment, and color scale changes help the model recognize essential features of the minority class, even though the number of original samples is limited [10]. This not only increases accuracy for the minority class but also helps reduce bias towards the majority class, making the model more efficient in identifying various leaf disease conditions [13, 14].

Traditional augmentation techniques such as image scaling, translation, or rotation are generally effective in improving DL image classifier accuracy. However, the impact of certain augmentations depends on the dataset's characteristics and the existing tasks [15]. Advanced/modern data augmentation strategies such as MixUp, CutOut, Mosaic, and Cutmix, which mix and remove different image regions, have recently gained popularity because the images produced by these approaches are still around the original data distribution [11, 12]. Models trained with this new data generalize well to data variations, avoid overfitting, and achieve robustness to data corruption by focusing on the complete structure of the object [16]. Advanced data augmentations like Generative Adversarial Network (GAN) introduce higher complexity and more substantial variation by creating new images that combine existing features [17].

This research draws inspiration from various image augmentation techniques previously developed, yet with a specific focus on addressing the issue of leaf disease classification in minority classes. In this literature review, the approach, EdgeCutMix, is compared with key other techniques discussed in the literature, such as MixUp [12], CutOut [18], CutMix [11], and Mosaic [19], highlighting the unique contribution to the field.

Below is an enhanced analysis of commonly used augmentation methods, highlighting their limitations and effectiveness, specifically in the context of leaf disease classification:

A. MixUp

MixUp [12] is a technique that generates new training samples by linearly interpolating between pairs of images and their labels. While it is praised for increasing model robustness and enhancing data diversity, MixUp often creates non-realistic images. This characteristic can lead to confusion in models, especially when distinguishing subtle features of leaf diseases, as the mixed images may not accurately represent valid disease presentations. Furthermore, studies have shown that while MixUp generally improves model generalization, it can impede performance in specialized tasks like leaf disease classification, where accurately identifying specific morphological features is crucial.

MixUp creates a new training sample ($\tilde{x}$, $\tilde{y}$) by combining two existing samples (x_i, y_i) and (x_j, y_j) using a weighted average:

$$\tilde{x} = \lambda x_i + (1 - \lambda) x_j \tag{1}$$

$$\tilde{y} = \lambda y_i + (1 - \lambda) y_j \tag{2}$$

where λ is a weight factor typically drawn from a Beta distribution Beta(α, α) for some $\alpha > 0$.

B. CutOut

CutOut [18] is a data augmentation technique that enhances model robustness by randomly masking out sections of input images during training. However, this method carries certain limitations; it might inadvertently remove critical features of diseases, particularly in instances where symptoms cover only a small area of the leaf, potentially leading to underfitting in minority classes. Despite these issues, CutOut can significantly improve robustness against overfitting. Nevertheless, the technique's inherent randomness can sometimes result in substantial information loss, which may negatively impact accuracy in specific disease classifications.

CutOut removes a random square patch from an input image (x), resulting in a modified image ($\tilde{x}$):

$$\tilde{x} = x \odot mask \tag{3}$$

Here, $\odot$ denotes element-wise multiplication, and mask is a binary matrix the same size as x where a randomly chosen region is set to zero.

C. Mosaic

Mosaic [19] is an augmentation technique that stitches together four training images into a single one, thereby enhancing the context variability that the model experiences during training. Although this method introduces high variability, it also increases the complexity of the learning task, which can potentially confuse the model regarding which features are relevant for disease classification. This technique offers better detection capabilities in complex scenes; however, its effectiveness in leaf disease classification may vary. This variation often depends on the degree of symptom overlap between the combined images, which can either aid or hinder the model's ability to correctly identify diseases.

Mosaic combines four training images (x_1, x_2, x_3, x_4) into a single new image ($\tilde{x}$) by arranging them in a 2×2 grid pattern:

$$\tilde{x}(r,c)=\begin{cases} x_1(r,c) \text{ if } r \le h/2 \text{ and } c \le w/2 \\ x_2(r,c) \text{ if } r \le h/2 \text{ and } c \le w/2a \\ x_3(r,c) \text{ if } r \le h/2 \text{ and } c \le w/2 \\ x_4(r,c) \text{ if } r \le h/2 \text{ and } c \le w/2 \end{cases} \tag{4}$$

where h and w are the image height and width, and (r,c) are pixel coordinates.

D. *CutMix*

CutMix [11] is an innovative data augmentation technique that merges elements of MixUp and CutOut by cutting and pasting patches between training images. Although this approach enhances data robustness, its random nature regarding patch selection and placement may disrupt the spatial coherence of disease signs. This disruption can lead to challenges in effectively learning important local features, potentially weakening the model's performance in specific contexts. While CutMix generally improves model robustness, it can be less effective for diseases characterized by localized symptoms. This is due to the technique's potential to obscure critical features necessary for accurate disease identification.

CutMix specifically cuts and pastes patches between training images. Given images (x_a, x_b), combining regions from both creates a new sample ($\tilde{x}$, $\tilde{y}$). A binary mask (M) defines the mixing ratio and area.

$$\tilde{x} = M \odot x_A + (1 - M) \odot x_B \tag{5}$$

$$\tilde{y} = \lambda y_A + (1 - \lambda) y_B \tag{6}$$

Here, $M \in \{0, 1\}$^(w x h) is the mask, λ is a mixing coefficient from Beta(α, α), and y_a, y_b are the one-hot encoded labels for x_A and x_B, respectively.

The region selection is based on a random bounding box $B = (r_x, r_y, r_w, r_h)$, where (r_x, r_y) is the top-left corner and (r_w, r_h) is the width and height.

CutMix encourages models to learn from broader features across different image parts, potentially improving robustness and generalization.

III. Materials and Methods

A. *Dataset Preparation*

The study used the Plant Pathology 2020 dataset from Kaggle's. It features 3,642 high-quality RGB apple leaf images (2,048×1,365 pixels) with annotations for diseases like scab, rust, multiple diseases, and healthy conditions, as detailed in Fig. 1. Most leaves show scab or rust symptoms, with only 5.0% displaying multiple diseases, indicating a significant imbalance. As found in several studies [20, 21], this dataset imbalance affects training, leading to higher accuracy for more common conditions over less frequent ones.

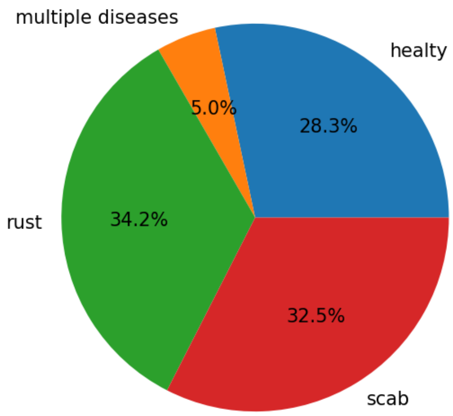

Fig. 1. Dataset distribution of the Plant Pathology 2020.

B. *EdgeCutMix Augmentation Technique*

The conceptual model titled "EdgeCutMix," depicted in Fig. 2, presents a workflow for processing a series of images, ranging from Image_1 to Image_n, to create an augmented image through structured steps. The process initiates with the random selection of 'm' images from the pool of available images. These selected images are then subjected to three primary stages of processing:

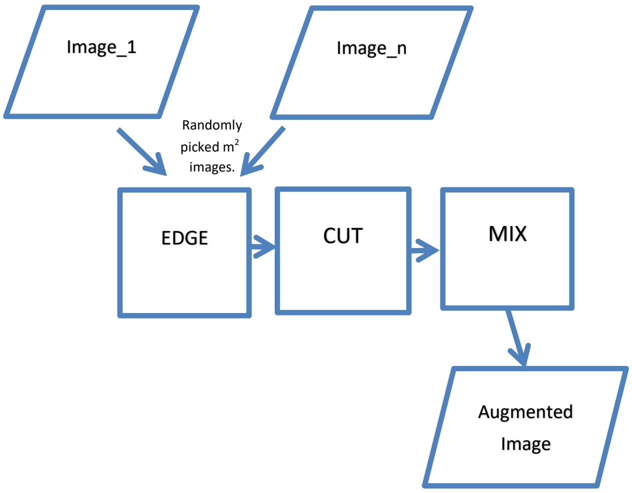

Fig. 2. Conceptual Model EdgeCutMix.

1) *Phase 1: EDGE*

The objective of this phase is to identify and highlight the critical edges within images that signify disease symptoms. For implementation, we applied the Canny edge detection algorithm to extract high-contrast edges, which serve as markers for important features in subsequent phases.

The first Phase involves applying an edge detection algorithm to identify image boundaries. The Canny edge detector, for instance, can be mathematically represented by the gradient of the image intensity function, where edge points are identified as local maxima of the gradient magnitude:

$$\nabla I = \sqrt{\left(\frac{\partial I}{\partial x}\right)^2 + \left(\frac{\partial I}{\partial y}\right)^2} \quad (7)$$

Here, I represent the image intensity function, and $\partial I/\partial x$ and $\partial I/\partial y$ are its partial derivatives along the x and y axes, respectively. The detected edges serve as a mask, highlighting the essential features for further processing.

2) Phase 2: CUT

The objective of this phase is to segment the images based on the detected edges, isolating important features for further augmentation. For implementation, we use the edge maps as masks to crop out Regions of Interest (ROI), preparing them for mixing. These ROIs focus on disease-specific features, ensuring their prominence in the augmented images.

Utilizing the mask M_e from the EDGE phase, the CUT phase focuses on cropping the regions of interest. The cropping is guided by a bounding box B calculated from M_e, which encapsulates the detected features. The operation can be visualized as applying the mask M_e to the original image I, then cropping to the bounding box:

$$I_c = I \odot M_e | B \quad (8)$$

where I_c is the cropped image, and $\odot$ denotes element-wise multiplication applied within the bounds of B.

3) Phase 3: MIX

The objective of this phase is to synthesize new images by intelligently blending cropped regions from multiple images. For implementation, cropped ROIs from different images are combined based on a set of predefined rules that maintain the natural appearance and variability of the diseases, resulting in a balanced and enriched dataset.

In the MIX phase, cropped images from different sources are combined into a single augmented image. This process involves creating a composite mask M_c that outlines how each cropped image Ic should be placed within the new image layout. Let $I_{a1}, I_{a2}, \ldots, I_{an}$ represent the combined augmented images, and $M_{c1}, M_{c2}, \ldots,$ *and* M_{cn} their corresponding masks. The final augmented image I_a can be represented as:

$$I_a = \sum_{i=0}^{n} I_{ai} \odot M_{ci} \quad (9)$$

Each $I_{ai} \odot M_{ci}$ represents a cropped image placed according to its mask within the larger composite image. The EdgeCutMix process can be seen in Fig. 2.

C. Augmented Target Dataset

The dataset of apple leaf diseases from Kaggle Plant Pathology 2020 consists of four classes: "Healthy", "Multiple_diseases", "Rust", and "Scab". The "Healthy" class has 516 images, while the "Multiple_diseases" class has the fewest samples, only 91 samples. The Rust and Scab classes have more samples, with 622 and 592 images, respectively. "Multiple_diseases" as the minority class poses a problem due to the lack of data samples. This leads to less accurate classification, especially because it shares characteristics with two other diseases, making it easily confused with other classes. To address this issue, data augmentation is performed on the "Multiple_diseases" class by adding 509 images, bringing the total to 600. This augmentation aims to balance the number of samples in each class and improve classification accuracy. Two randomly selected images from the 'Rust' and 'Scab' classes, each with dimensions of 512×512 pixels, are arranged in a 2×2 grid. Each grid cell represents an image of the same size, 512×512 pixels. The images from the "Multiple_diseases" class that have been augmented to reach 509 images are then resized to fit the grid size (512×512 pixels).

D. Experimental Setup

The implementation was carried out using Python, incorporating the OpenCV library for image processing and TensorFlow for model training. The selection and tuning of hyperparameters were pivotal in optimizing the performance of the model. Guided by preliminary tests that evaluated sensitivity to various settings, specific parameters such as the learning rate, number of epochs, and batch size were meticulously optimized. Initially, a learning rate of 1e–4 was chosen, which was adjusted based on the plateauing of validation loss observed during initial trials. The number of epochs was set at 25 after it was noted that prolonged training did not significantly enhance performance and instead increased the risk of overfitting. The batch size was established at 32 to strike a balance between computational efficiency and the stability of gradient updates.

E. Validation of Augmented Data

EdgeCutMix augmented data was validated using a feature distribution analysis approach. This entailed employing statistical methods to analyze and compare the feature distributions between the original and augmented datasets, focusing on measures such as mean, median, mode, skewness, and kurtosis of critical features. The goal was to verify that the augmentation process preserved the essential characteristics of the data while introducing desirable diversity. This analysis ensures that any changes in the statistical distribution of features do not compromise the data's integrity [22].

F. Model Training and Evaluation

EdgeCutMix was implemented in Python using OpenCV and TensorFlow libraries. Four CNN architectures were used to evaluate the effectiveness of EdgeCutMix: MobileNetV2, EfficientNetB7, ResNet50, and DenseNet201. Each model was initialized with weights pre-trained on the ImageNet dataset. Training was conducted using a split of 80% of the images for training and 20% for validation. The Adam optimizer, with a learning rate 1e–4, was utilized, and models were trained for 25 epochs. The training process incorporated standard data augmentation techniques (rotation, zoom, and horizontal flipping) alongside EdgeCutMix to investigate the synergistic effect on model performance. During the training process, fine-tuning was performed by adjusting the learning rate to monitor overfitting.

G. Performance Evaluation

The performance evaluation was conducted based on accuracy derived from the confusion matrix to assess the model's effectiveness across all classes comprehensively. Additionally, loss was continuously monitored to detect any signs of overfitting.

IV. Result and Discussion

A. Augmented Data

The following figure presents a comparative view of innovative data augmentation methods to enhance model robustness and improve learning efficiency. Fig. 3 illustrates examples of five distinct data augmentation techniques. The CutOut technique is demonstrated by a randomly placed black box on a leaf, eliminating some visual information as a form of regularization. MixUp depicts two leaf images superimposed on each other with transparency, amalgamating the visual features and labels from both images to generate a more diverse input for the model. CutMix appears as a segment from one leaf placed on another, creating a hybrid image with parts from both sources.

Meanwhile, Mosaic combines four leaf images into one by dividing the image area into four segments, each displaying a portion of a different image. This allows the model to learn from a more complex context within a single image. Lastly, Cutegdmix presents multiple leaves with cuts that appear smoother and more integrated, indicating that this technique prioritizes precision in cutting and more natural integration of the mixed images. Mosaic and Cutegdmix offer more significant variation within a single image than the other techniques.

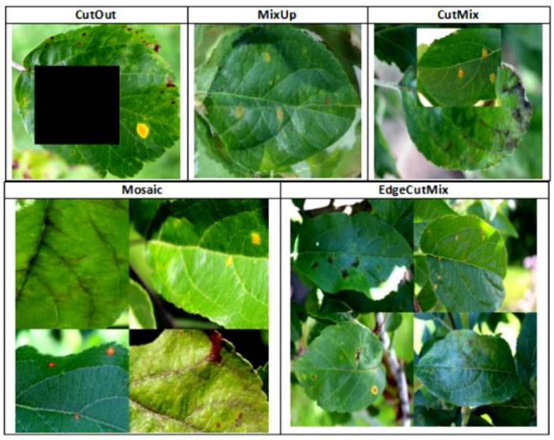

Fig. 3. Visualize a sample image of mixing augmentation.

B. Validation of Augmented Data

Validating augmentation in this way helps ensure that augmented data maintains the quality and essential characteristics of the original data while adding to its diversity. This is crucial for building robust deep-learning models that generalize well to new data.

The validation of data resulting from augmentation The statistical analysis employed to validate the augmented data focuses on metrics such as mean, median, and standard deviation, which are crucial for assessing the preservation of essential characteristics in augmented images. These metrics help ensure that the augmented data maintains a distribution similar to the original dataset, thus preserving the critical disease indicators necessary for accurate classification. This process reassures us that the augmentation has not introduced any bias or distortion that could mislead the training of the model.

The validation of data resulting from augmentation is conducted using statistical analysis. This statistical analysis employs minimum, maximum, average, median, and range metrics to identify disease patterns in leaves for each disease class. Visualization uses a histogram of RGB color distribution complemented by a Boxplot diagram.

Fig. 4 shows that for the category of multiple diseases, the intensity of color tends to be higher in the green channel than in the red and blue channels. The variability of pixel intensity (as indicated by the standard deviation) is highest in the green channel. The pixel intensity distribution in the red and blue channels tends to be more symmetrical around the median, considering that the median and mean are very close.

The Red Box Plot: With a median value of 95 and a range from 47 to 127, it is observed that the median in the red box plot is precisely positioned in the middle of the box, and the whiskers encompass the range from 47 to 127. This is consistent with the previous interpretation of the red box plot. The Green Box Plot: With a median value of 129 and a range from 85 to 165, this should also be reflected in the green box plot, with the median situated in the middle and the whiskers encompassing the range above, by the previous description of the green box plot. The Blue Box Plot: With a median value of 67 and a range from 28 to 106, it is observed that the median in the blue box plot is closer to the lower quartile of the box. The whiskers extend from 28 to 106, consistent with the previous description of the blue box plot.

Based on the statistical results of the augmented data from EdgeCutMix Table I, the values within the table remain within the range of values on the box plot in Fig. 4. Table I shows the statistical results of the augmented data from EdgeCutMix. The values within the table remain within the range of values on the box plot in Fig. 4, indicating that the augmentation process preserves the essential characteristics of the data while introducing desirable diversity. This validation is crucial for ensuring that any changes in the statistical distribution of features do not compromise the data's integrity.

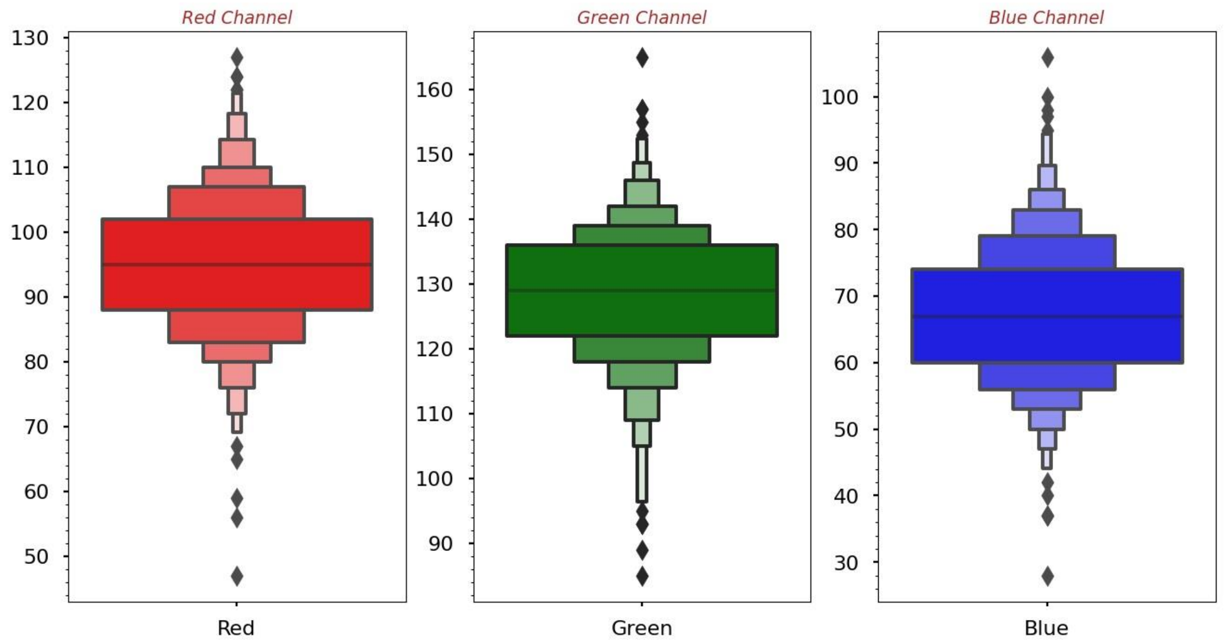

Fig. 4. The values of Red, Green, and Blue (RGB) channels are associated with multiple diseases.

TABLE I. THE EDGECUTMIX STATISTICS OF RGB CHANNELS

Channel	MIN	MAX	MEAN	MEDIAN	RANGE	STDDEV
RED	47	127	95.1	95	80	10.417939
GREEN	85	165	128.7	129	80	10.168322
BLUE	28	106	67.3	67	78	10.163791

C. Training the Classification Model

The CNN model is trained using the training data to recognize patterns and features associated with each type of apple leaf disease. The dataset is divided into two parts, namely training data (80%) and validation data (20%) to train and test the model. The CNN model with MobileNetV2, EfficientNetB7, ResNet50, and DenseNet201 architectures is used to identify apple leaf diseases based on the provided images.

Fig. 5 compares the performance of four different deep learning model architectures using the EdgeCutMix data augmentation technique. There are four models: MobileNetV2, EfficientNetB7, ResNet50, and DenseNet201. Based on Fig. 5, the EfficientNetB7 model with EdgeCutMix obtained the best results regarding accuracy and validation loss, indicating that this architecture is more suitable for this type of data and classifying apple leaf diseases than other architectures in this experiment.

The EdgeCutMix technique on the EfficientNetB7 model performance showed a perfect training accuracy of 100% with a high validation accuracy of 98 %.

Training our classification model provided several insights and also presented challenges. For instance, we encountered convergence issues in early training phases, which were mitigated by adjusting learning rates and optimizing batch sizes. Furthermore, the selection and tuning of hyperparameters, particularly in balancing the trade-off between model complexity and performance, were challenging. These experiences underscore the delicate balance required in deep learning setups and highlight the importance of robust experimental design.

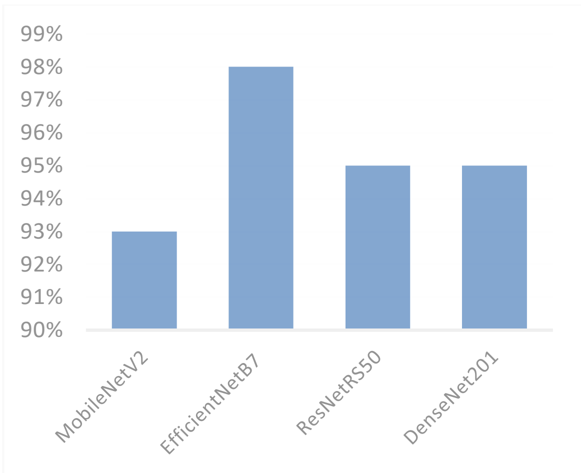

Fig. 5. Validation accuracy of EdgeCutMix.

D. Fine-Tuning

During the training process of the CNN model, it is crucial to monitor overfitting by examining the Training Loss vs Validation Loss graph. If the Training Loss continues to decrease while the Validation Loss begins to increase, this indicates overfitting. To overcome overfitting, fine-tuning can be done by adjusting the learning rate. The CNN model can become more generalized and effectively classify apple leaf diseases by fine-tuning.

The fine-tuning performed on classifying apple leaf diseases involves using a learning rate strategy. This strategy aims to take advantage of the initial Phase of training where the model can learn quickly (ramp-up), then allow the model to consolidate what it has learned at a higher learning rate (sustain), and finally minimize the learning rate to allow the model to make fine adjustments to its weights. The implementation of the "learning rate warm-up" strategy where the learning rate gradually increases at the beginning of training to aid model convergence, then is kept constant to give the model time to find a good local minimum, and finally reduced to allow the model to refine adjustments at that minimum. The learning rate increases significantly at the beginning of training (around epoch 0 to 20). Then, the learning rate is kept constant from epoch 20 to around 80. Finally, the learning rate decreases sharply to approach 0 at epoch 80.

Based on the training results of the CNN model with fine-tuning, the model accuracy graph and model loss graph in Fig. 6 can be analyzed as follows:

1. Initial Training Phase: Both graphs show a significant increase in accuracy and a decrease in loss for the training and testing sets in the initial training phase. This indicates that the initial increasing learning rate strategy (ramp-up) has been successful in helping the model to learn quickly.

2. Sustain Phase: After the ramp-up phase, the learning rate remains constant for several epochs (sustain Phase). During this period, the model's accuracy increases slowly, and loss decreases, indicating that the model has achieved a balance between learning new patterns and reinforcing previously learned patterns without significant overfitting.

3. After the Sustain Phase: After the ramp-up and sustain phases, the learning rate should decrease, but based on the graphs, no performance decline indicates the need for a significant decrease in the learning rate. Thus, maintaining a low learning rate (after ramp-up and sustain) has been adequate to maintain model convergence.

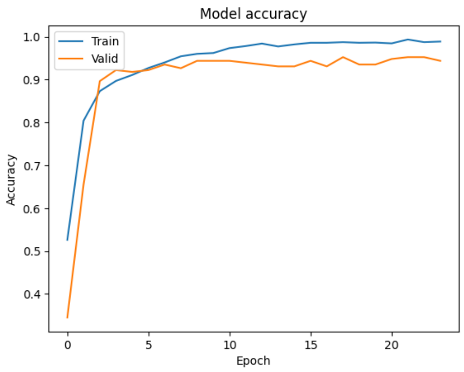

(a)

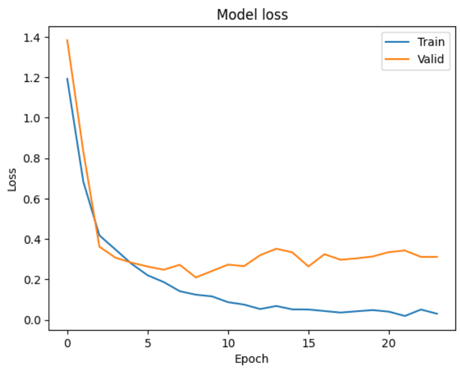

(b)

Fig. 6. Training results (a) accuracy (b) loss.

The accuracy of the testing set remains stable above 90%, and the loss remains stable below 0.4, indicating excellent performance from the model. There are no apparent signs of overfitting, as the testing accuracy does not decrease, and the testing loss does not show a consistent increase.

The learning rate strategy has successfully achieved high performance without causing significant overfitting. The ramp-up phase accelerates initial learning, the sustain phase allows for learning consolidation, and the subsequent learning rate adjustments successfully maintain good performance.

Overall, the results obtained from the effective learning rate strategy are in line with the trends shown in the graph. The model experiences a steady increase in accuracy and a consistent decrease in loss throughout the training process, indicating that the learning rate strategy contributes positively to the model's convergence and generalization.

Table II compares the performance of the CNN model using the "EdgeCutMix" augmentation technique with models using the "MixUp," "CutOut," "CutMix," and "Mosaic" augmentation techniques. The results highlight that EdgeCutMix combined with base augmentation methods provides a good balance between training and validation accuracy, reducing overfitting while maintaining high accuracy. This comparison demonstrates the effectiveness of EdgeCutMix in managing minority classes and improving overall model performance.

TABLE II. COMPARISON OF TRAINING AND VALIDATION VALUES

No	Method	Training Accuracy	Training Loss	Validation Accuracy	Validation Loss
1	MixUp	1.0000	0.0004	0.9506	0.2867
2	MixUp+AugBase	0.9705	0.0756	0.9464	0.2770
3	CutOut	1.0000	0.0004	0.9356	0.5045
4	CutOut+AugBase	0.9898	0.0330	0.9700	0.1513
5	CutMix	1.0000	0.0003	0.9506	0.3856
6	CutMix+AugBase	0.9775	0.0761	0.9506	0.1655
7	EdgeCutMix	1.0000	0.0001	0.9700	0.1704
8	EdgeCutMix+AugBase	0.9855	0.0494	0.9678	0.1194
9	Mosaic	0.9989	0.0026	0.9657	0.2051
10	Mosaic+AugBase	0.9871	0.0427	0.9528	0.2024

From Table II, several points can be derived the "MixUp," "CutOut," and "CutMix" methods, when used alone without base augmentation, provide perfect accuracy on training data but not on validation data. This indicates overfitting, where the model fits too well with the training data and is less generalizable to new data.

When base augmentation is combined with other methods, there is a slight decrease in training accuracy but an increase in validation accuracy, indicating a more general model and better generalization on unseen data.

In general, the loss values for the training data are shallow, which indicates that the model learns well from the training data. However, the higher loss values in the validation data indicate that there is room for improvement in terms of model generalization.

The "CutMix+AugBase" and "EdgeCutMix +AugBase" methods provide a good balance between training and validation accuracy compared to other methods; this shows that these two methods combined with base augmentation can reduce overfitting while still maintaining high accuracy.

A significant discrepancy between training and validation accuracy, or between training loss and validation loss, may indicate issues such as overfitting. Perfect training accuracy coupled with lower validation accuracy suggests the potential for overfitting.

The analysis of experimental results provided significant insights into the comparative performance of different augmentation techniques. Notably, EdgeCutMix demonstrated superior efficacy over traditional methods like MixUp and CutOut, especially in managing minority classes. The selective mixing strategy employed by EdgeCutMix, which focuses on preserving critical features within images, was instrumental in its enhanced performance. These findings highlight the potential importance of edge detection in refining augmentation techniques for image-based classifications. Looking forward, there is a promising avenue for future research to integrate more complex edge detection algorithms or to combine EdgeCutMix with other advanced data augmentation forms to further bolster model robustness.

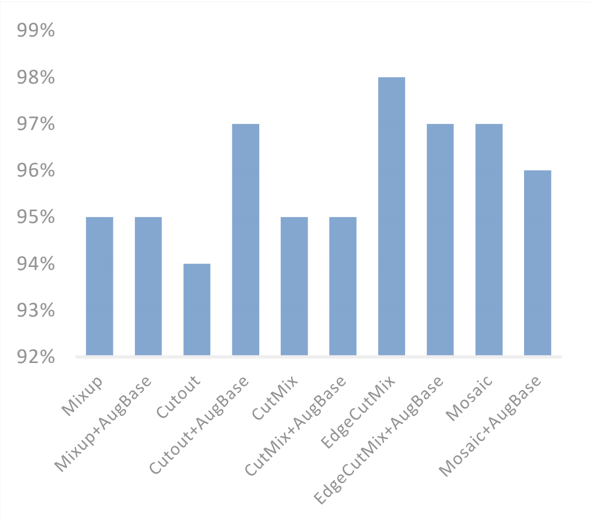

Fig. 7. Comparison of accuracy value based on the confusion matrix.

Fig. 7 shows the accuracy values of various image augmentation techniques calculated based on the confusion matrix for an apple leaf disease classification model. The accuracy value here represents the total proportion of correct predictions to the total number of predictions made by the model. These techniques produce relatively high accuracy, with EdgeCutMix showing the best performance at 98%. From the experiments, it was found that basic augmentation sometimes increases and sometimes decreases accuracy, depending on the augmentation technique used. This indicates that augmentation techniques should be used carefully to achieve the best results on a specific dataset.

E. Discussion

The augmentation method of blending operates by randomly selecting two or more images from the training data. Cutegdmix is a mixed augmentation method encompassing digital image processing, bounding box information storage, and creating puzzles from these images.

Table III compares the Cutegdmix method with four existing methods based on five criteria. The comparison highlights that Cutegdmix excels in cropping precision, which is crucial for retaining disease-specific features. This table helps in understanding the strengths and limitations of different augmentation techniques, guiding the selection of appropriate methods based on specific dataset characteristics and desired outcomes.

MixUp, CutMix, Mosaic, and Cutegdmix utilize the entire image region and blend images and labels. However, only Cutegdmix has high cropping precision, indicating more accurate and appropriate cropping in line with the essential features in the image. CutOut does not blend images or labels and only removes a specific part of the image. In terms of the number of image mixings, MixUp and CutMix blend two images, while Mosaic and Cutegdmix blend four images, suggesting that the latter two techniques provide richer data variation for training machine learning models.

TABLE III. COMPARISON AMONG MIXING AUGMENTATION METHODS

Comparison	MixUp	CutOut	Cutmix	Mosaic	EdgeCutmix
Usage of the full image region	√	×	√	√	√
Regional dropout	×	√	√	√	√
Mixed image and label	√	×	√	√	√
Precision of Cropping	×	×	×	×	√
Number of image mixings	2	1	2	4	4

In Table III, while Cutegdmix shows promise, its qualitative comparison with methods like MixUp and Mosaic highlights that no single method is universally superior. Cutegdmix excels in precision of cropping, which is crucial for retaining disease-specific features. However, it may fall short in scenarios requiring broader contextual understanding where methods like Mosaic might perform better. Understanding these nuances helps in selecting the appropriate augmentation technique based on specific dataset characteristics and desired outcomes.

V. CONCLUSION

In this study, we introduced EdgeCutMix, an innovative image augmentation technique specifically designed to address class imbalances in leaf disease classification. Through the integration of edge detection and selective mixing, EdgeCutMix has demonstrated significant improvements in classification accuracy, particularly for minority classes which are often underrepresented in training datasets. This approach involves generating additional samples for minority classes, ensuring a more balanced representation within the training dataset.

EdgeCutMix contributes significantly to the advancement of data augmentation techniques by offering a method that not only enhances the diversity and balance of training datasets but also preserves critical information essential for accurate classification. This approach addresses one of the key challenges in deep learning, promoting more robust and generalizable models. By pushing the boundaries of traditional augmentation methods, EdgeCutMix sets a new standard for how we can creatively leverage image processing techniques to improve deep learning outcomes.

Based on the experiment, EdgeCutMix significantly improves classification accuracy for minority classes, achieving up to 98% accuracy. EdgeCutMix represents a substantial step forward in tackling the prevalent issue of class imbalance in image classification. Its development underscores the potential of thoughtful, targeted interventions in data processing to yield significant improvements in model performance across various domains. This study not only enhances our understanding of augmentation techniques but also opens up new pathways for leveraging these strategies to solve real-world problems effectively.

EdgeCutMix's utility extends beyond the realm of agricultural pathology. Its application could revolutionize image-based classification in fields such as surveillance or autonomous driving, which might involve addressing specific challenges like varying lighting conditions, motion blur, or occlusions. Future research could explore several avenues to enhance the effectiveness and applicability of EdgeCutMix. One promising direction is the integration with advanced deep learning strategies, such as transfer learning or Generative Adversarial Networks (GANs), to further improve model performance. Another potential area of exploration is optimizing EdgeCutMix for real-time applications, like in-field disease detection using mobile devices, making the technology more practical and accessible.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Derisma was responsible for carrying out the research work and took the lead in writing the manuscript under the guidance and supervision of Nur Rokhman and Dyah Aruming Tyas. All authors approved the final version.

References

[1] R. M. Math and N. V. Dharwadkar, "Early detection and identification of grape diseases using convolutional neural networks," *Journal of Plant Diseases and Protection*, vol. 129, no. 3, pp. 521–532, 2022.

[2] K. D. Priyadharsini and B. Kalpana, "Survey on disease detection in plant leaf using data augmentation," in *Proc. 2021 5th International Conference on Electronics, Communication and Aerospace Technology (ICECA)*, 2021, pp. 1686–1690.

[3] J. G. A. Barbedo, "Identification of plant diseases from individual lesions and spots utilizing deep learning," *Biosystems Engineering*, vol. 180, pp. 96–107, 2019.

[4] J. Lu, L. Tan, and H. Jiang, "Review on Convolutional Neural Network (CNN) applied to plant leaf disease classification," *Agriculture*, vol. 11, no. 8, 707, 2021. doi: 10.3390/agriculture11080707

[5] A. R. Khan *et al.*, "Automated diagnosis of rice leaves diseases using a hybrid deep learning model," *Journal of Advances in Information Technology*, vol. 14, no. 3, pp. 418–425, 2023. doi: 10.1109/icaiic57133.2023.10067113

[6] M. Arsenovic *et al.*, "Solving current limitations of deep learning based approaches for plant disease detection," *Symmetry*, vol. 11, no. 7, 939, 2019. doi: 10.3390/sym11070939

[7] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, pp. 1263–1284, 2009. doi: 10.1109/TKDE.2008.239

[8] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, 2019.

[9] S. M. Hassan *et al.*, "Identification of plant-leaf diseases using CNN and transfer-learning approach," *Electronics*, vol. 10, no. 12, 1388, 2021. doi: 10.3390/electronics10121388

[10] S. Yadav and G. P. Bhole, "Handling imbalanced dataset classification in machine learning," in *Proc. 2020 IEEE Pune Section International Conference (PuneCon)*, Dec. 2020, pp. 38–43.

[11] S. Yun *et al.*, "CutMix: Regularization strategy to train strong classifiers with localizable features," in *Proc. 2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 6022–6031.

[12] H. Zhang *et al.*, "MixUp: Beyond empirical risk minimization," arXiv Preprint, arXiv: 1710.09412, 2017.

[13] P. Chen *et al.*, "GridMask data augmentation," arXiv Preprint, arXiv: 2001.04086, 2020.

[14] E. D. Cubuk *et al.*, "Randaugment: Practical automated data augmentation with a reduced search space," in *Proc. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020. doi: 10.1109/CVPRW50498.2020.00359

[15] E. D. Cubuk *et al.*, "AutoAugment: Learning augmentation strategies from data," in *Proc. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. doi 10.1109/CVPR.2019.00020

[16] Z. Zhong *et al.*, "Random erasing data augmentation," arXiv Preprint, arXiv:1708.04896, 2017.

[17] Derisma, N. Rokhman, and I. Usuman, "Systematic review of the early detection and classification of plant diseases using deep learning," in *Proc. IOP Conference Series: Earth and Environmental science, 1st Lekantara Annual Conference on Natural Science and Environment*, Malang, 2022, vol. 1097. doi: 10.1088/1755-1315/1097/1/012042

[18] T. DeVries and G. W. Taylor, "Improved regularization of convolutional neural networks with CutOut," arXiv Preprint, arXiv: 1708.04552, 2017.

[19] A. Bochkovskiy, C. Wang, and H. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," arXiv Preprint, arXiv: 2004.10934, 2020.

[20] Liu *et al.*, "Identification of apple leaf diseases based on deep convolutional neural networks," *Symmetry*, vol. 10, no. 1, 11, 2018. doi: 10.3390/sym10010011

[21] A. Yadav *et al.*, "AFD-Net: Apple foliar disease multi-classification using deep learning on plant pathology dataset," *Plant and Soil*, vol. 477, no. 1, pp. 595–611, 2022.

[22] J. Jung and K. W. Kim, "Study of validation methods for augmented data," in *Proc. 2023 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)*, 2023. doi: 10.1109/icaiic57133.2023.10067113