

Lacunarity Pooling Layers for Plant Image Classification using Texture Analysis

Akshatha Mohan and Joshua Peebles

Department of Electrical and Computer Engineering, Texas A&M University
College Station, TX, USA

{akshatha.mohan, jpeeples}@tamu.edu

Abstract

Pooling layers (e.g., max and average) may overlook important information encoded in the spatial arrangement of pixel intensity and/or feature values. We propose a novel lacunarity pooling layer that aims to capture the spatial heterogeneity of the feature maps by evaluating the variability within local windows. The layer operates at multiple scales, allowing the network to adaptively learn hierarchical features. The lacunarity pooling layer can be seamlessly integrated into any artificial neural network architecture. Experimental results demonstrate the layer's effectiveness in capturing intricate spatial patterns, leading to improved feature extraction capabilities. The proposed approach holds promise in various domains, especially in agricultural image analysis tasks. This work contributes to the evolving landscape of artificial neural network architectures by introducing a novel pooling layer that enriches the representation of spatial features. Our code is publicly available.¹

1. Introduction

In computer vision, the common assumption of uniform pixel intensity in images fails to capture the complexity inherent in real-world images [7]. Visual textures that are easily understood by humans introduce intricate patterns, because of surface roughness and/or color disparities [2]. These natural surfaces exhibit spatial distribution of local image patterns can exhibit statistical self-similarities. The notion of self-similarity is that an object is composed of sub-units and sub-sub-units on multiple levels that (statistically) resemble the structure of the whole object [15].

For Agricultural image analysis, understanding the intricate textures and spatial patterns of plant images plays a pivotal role in tasks such as disease detection, crop monitoring, and yield prediction [9]. Methods that focus on color

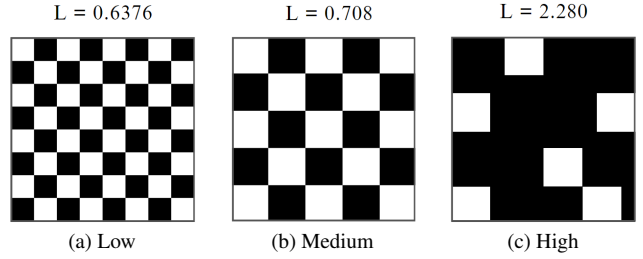


Figure 1. Fractal patterns exhibit an increase in lacunarity from left to right, indicating a rise in irregular gaps towards the right.

and shape features may overlook the rich information encoded in image textures [11]. Therefore, there is a need to introduce novel techniques to capture the spatial heterogeneity of plant images. Lacunarity, a measure traditionally employed in fractal analysis, offers a unique perspective on texture analysis by quantifying the variability of pixel intensities within local windows or globally as shown in Figure 1. We propose a novel texture analysis approach grounded in the concept of lacunarity [19] to address the limitations of existing methods.

We propose a lacunarity pooling layer as an alternative to the conventional pooling layers. This layer leverages the spatial distribution of gap or hole sizes within images, providing a scale-dependent measure of texture. By integrating lacunarity into the network's feature extraction process, we aim to capture intricate spatial relationships for effective image classification. Our work presents several key contributions. Firstly, we introduce the lacunarity pooling layer, a novel addition to neural network architectures, which incorporates multiple methods for calculating lacunarity. Secondly, we propose several different methods to calculate lacunarity include gliding box, differential box counting, and multi-scale approaches. All of these approaches can be implemented using pre-existing layers (e.g., average and sum pooling). Lastly, the implementation of the lacunarity pooling layer allows for the extraction of intricate spatial features from input data, enabling more robust and versatile representations to be learned by the network.

¹https://github.com/Advanced-Vision-and-Learning-Lab/2024_V4A_Lacunarity_Pooling_Layer




Figure 2. Overview of the proposed fusion model. Images are passed into a frozen pre-trained network and the feature maps from the final convolution layers are passed into the lacunarity pooling layer. The lacunarity features maps are then multiplied by the feature maps from the average pooling layer for enhanced analysis.

2. Related Work

Lacunarity has been integrated with convolutional neural networks (CNNs) to improve feature extraction and enhance accuracy in various applications. The lacunarity cue with CNN is a novel method for saliency detection in synthetic aperture radar (SAR) images [18]. Lacunarity identifies regions of interest (ROIs) by measuring the irregularity of scattering points within superpixels. By leveraging the irregular fluctuations in backscattering intensity, lacunarity effectively highlights ROIs in SAR imagery [18]. Additionally, lacunarity has successfully characterized seafloor textures in synthetic aperture sonar (SAS) imagery [28].

A related approach was formulated by Florindo [10] where a novel technique called fractal pooling is introduced. The conventional CNN approach for texture analysis using fully-connected layers for classification often falls short in achieving optimal performance due to the lack of invariance to rigid transforms like translation, rotation, and scaling. Global average pooling (GAP) has been traditionally employed to address this issue, but GAP results in the loss of spatial information and relationships across multiple scales. To overcome this challenge, fractal dimension is explored as an alternative pooling method for deep convolutional features in texture representation. Fractal dimension is computed over the last convolutional layer using CNN modules, summarizing higher-level features and offering a promising approach to encoding and pooling deep texture features effectively.

Different fractal sets might have identical fractal dimensions while exhibiting diverse appearances or textures [16]. Mandelbrot introduced the term lacunarity to describe that characteristic of fractals of the same dimension with different appearances or textures [20]. Lacunarity serves as the

complementary metric to fractal dimension, providing insight into the spatial arrangement of gaps or holes within an object. The variability or texture in gap sizes within the spatial structure of geometric objects determines the level of lacunarity. Fractal dimension measures complexity, but can give a constant value across different visual structures. Lacunarity addresses this by quantifying the “gappiness” or heterogeneity in images, providing a more nuanced insight into their structural characteristics.

Deep Texture Recognition via Exploiting Cross-Layer Statistical Self-Similarity [5] introduces a novel approach to texture recognition by leveraging the multi-scale nature of texture features and the self-similarity properties of textures. By analyzing the evolution of texture structures across different layers of a convolutional neural network, the proposed CLASSNet model captures multi-scale information inherent in textures. Additionally, the model exploits statistical self-similarity, a well-known property of textures characterized by fractal dimensions, to enhance the discriminative power of feature aggregation. This integration of multi-scale analysis and self-similarity fractals in texture recognition enables CLASSNet to effectively encode texture information at varying scales and improve classification accuracy. Our proposed lacunarity pooling layer also aims to incorporate multi-scale feature analysis to effectively capture texture information.

3. Methodology

3.1. Lacunarity Pooling Model

We introduce a new model, designed for integrating pre-trained CNN architectures with additional lacunarity pooling layers for improved feature extraction and classification tasks. As seen in Figure 2, the model takes an input image

and extracts features using any pre-trained backbone (*i.e.*, ConvNexT, ResNet18). For our work, the pre-trained backbone is fixed and only the final output classification layer is tuned. The backbone is fixed 1) to reduce computational costs of fine-tuning the model end-to-end and 2) to evaluate if the lacunarity pooling layer can capture features from the backbone network that are effective for classification.

The pooling layer, determined by the specified configuration, is used to spatially aggregate the feature maps. Following inspiration from [10], a global average pooling (GAP) layer is employed to capture global context from the feature maps. The fusion process combines the lacunarity feature maps with the GAP feature maps, enhancing the model’s ability to capture both local and global spatial information. Finally, the fused features are fed into a fully connected layer for making predictions. This flexible model architecture provides a robust framework for leveraging both pre-trained CNNs and lacunarity-based features. We investigated three approaches to compute lacunarity and each calculation is discussed in the next three subsections.

3.2. Base Lacunarity

Lacunarity is based on the pixel distribution for an image, which is obtained from scans for different box sizes at various grid orientations. Lacunarity [28] can be defined through Equation 1:

$$L = \frac{\sigma^2}{\mu^2} \quad (1)$$

where σ^2 (variance) = $\frac{1}{n} \sum_{i \in N} (x_i - \mu)^2$ and μ (mean) = $\frac{1}{n} \sum_{i \in N} (x_i)$. This formulation quantifies the heterogeneity or clustering of pixel intensities, providing a measure of lacunarity for texture analysis.

L can be therefore rewritten as shown in Equation 2:

$$L = \frac{n \sum_{i=1}^n x_i^2}{(\sum_{i=1}^n x_i)^2} - 1 \quad (2)$$

Our proposed base lacunarity pooling can be easily implemented in any deep learning framework. The numerator can simply be computed by using a sum pooling layer applied to the squared features maps multiplied by the size of the kernel (n). Similarly, the denominator squares the sum pooled feature maps.

By examining the distribution of gaps or holes at various scales, lacunarity provides valuable insights into the structural complexity and heterogeneity of the image. The process of computing lacunarity involves applying normalization to the input through the hyperbolic tangent (tanh) function to ensure that the feature values are transformed to a range between -1 and 1. To adapt the feature maps to typical pixel values ranging from 0 to 255, an additional transformation is applied as shown in Equation 3:

$$x = \left(\frac{\tanh(x) + 1}{2} \right) \times 255 \quad (3)$$

This equation scales the output of the tanh function to the desired range for pixel values (0 to 255) commonly used in image processing. The combination of these two formulas allows for effective normalization of the input data.

3.3. Differential Box Counting Lacunarity

The differential box counting (DBC) method is another approach for lacunarity computation [8]. The box counting algorithm is only valid for binary images, and therefore, it is not applicable to grayscale images, as required by DBC algorithm. Briefly, DBC algorithm determines the fractal dimension value of a grayscale image by performing a least square linear fit of $\log(N_r)$ with respect to $\log(1/r)$, where N_r represents the box-counting for a grid of size r pixels. The box-counting process for a grid of size r involves computing the difference between the boxes that denote the maximum and minimum gray levels of the pixels within the grid.

The DBC module is a novel approach to estimate lacunarity. Given input image(s) or feature map(s), the module normalizes pixel intensities and performs dilated box counting through max and min pooling operations. The lacunarity at different dilation factors r is computed using the formula shown in Equation 4:

$$L_r = \frac{(M_r^2) \cdot Q_{mr}}{(M_r \cdot Q_{mr} + \epsilon)^2} \quad (4)$$

Here, M_r is the sum of n_r , representing the count of boxes of size r containing pixels, and Q_{mr} is the average occupancy of these boxes. The results for various r values are concatenated and processed through a convolutional layer, enabling the module to capture heterogeneity in pixel intensities within specific regions of the image.

In the context of analyzing mass distribution in sets, Alain and Cloitre [1] introduced a gliding-box algorithm to define lacunarity. This method involves a box of radius r sliding across a lattice superimposed on the set. Here, $n(M, r)$ represents the count of gliding-boxes with radius r and mass M . The probability function $Q(M, r)$ is derived by normalizing $n(M, r)$ by the total box count.

The mass M in Equation 4 can be computed using a DBC method introduced by Sarkar and Chaudhuri [26]. As shown in Figure 3, a cubic of size $r \times r \times r$ is positioned over the upper-left corner of an image window with dimensions $W \times W$, where W is an odd number, and $r < W$. Based on pixel values, a column with more than one cubic may be required to cover the image intensity surface. Each r -th gliding-box is assigned numbers (*e.g.*, 1, 2, 3) from bottom to top. For a given r -th gliding-box, with minimum and




Figure 3. Illustration of differential box counting. The input image is divided into window sizes $w \times w$, and further subdivided into $r \times r \times r$ sized cubes stacked up until it covers the grayscale intensity of the local window. Here, u represents the lowest grayscale value and v represents the highest grayscale value.

maximum pixel values falling in boxes u and v respectively, the relative height of the column is given by Equation 5:

$$n_r(i, j) = v - u + 1 \quad (5)$$

where i and j denote spatial coordinates. v and u can be computed using max and min pooling. As the r -th gliding-box traverses the $W \times W$ image window, the total mass is calculated through Equation 6:

$$M_r = \sum_{i,j} n_r(i, j) \quad (6)$$

3.4. Multi-scale Lacunarity

The multi-scale method involves computing lacunarity from different resolution levels in a Gaussian pyramid [3] as shown in Figure 4. The image undergoes recursive down-sampling and convolution with a Gaussian filter. The Kornia library [25] is used to construct the gaussian pyramid. At each level of the pyramid, lacunarity is calculated using Equation 2, capturing the spatial distribution of feature variations. Subsequently, these lacunarity values are upsampled using the bilinear interpolation to the original resolution, providing a comprehensive representation of lacunarity at different scales. The concatenated lacunarity values are then processed through a convolutional layer.

The lacunarity values are calculated at multiple scales for each channel resulting in CS features maps where C is the number of input channels and S is the number of scales. We use a 1×1 convolution operation to perform to learn a linear combination of the different scales for each channel to ensure the number of input and output channels are the same. This will allow the layer to be easily integrated into any existing network to replace other pooling layers. This approach allows for a multi-scale analysis, enabling the model to capture intricate features at various levels of granularity within the image. The multi-scale lacunarity pooling layer is summarized in Algorithm 1.




Figure 4. Multi-scale lacunarity feature approach is shown. Each input feature map is computed at different scales using a Gaussian pyramid. For each scale, the lacunarity feature map is computed and each scale is then upsampled to the same spatial resolution before concatenation.

Algorithm 1 Multi-scale Lacunarity Pooling Operation

Require: $\mathbf{X} \in \mathbb{R}^{N \times C \times H \times W}$ feature maps, $f_\theta(\mathbf{X})$ multi-scale operation (e.g., Gaussian pyramid), scale factor S

- 1: Get scale pyramid feature maps, $\mathbf{Z} \leftarrow f_\theta(\mathbf{X})$
- 2: **for** $s = 1$ to S **do**
- 3: Select scale s feature maps, \mathbf{Z}_s
 Compute lacunarity feature maps, \mathbf{L}_s , using Equation 2
 Upsample spatial dimension of feature maps based on \mathbf{L}_1 , $H' \times W'$
 Concatenate lacunarity \mathbf{L}_s on channel dimension
- 4: **end for**
- 5: Use 1×1 convolution to reduce $\mathbf{L} \in \mathbb{R}^{N \times CS \times H' \times W'}$ to $\mathbf{L} \in \mathbb{R}^{N \times C \times H' \times W'}$ to combine lacunarity across different scales
- 6: **return** \mathbf{L}

4. Experimental Results

4.1. Experimental Setup

The following experimental setup was used for training: a batch size of 128, cross-entropy loss function, Adam optimization with an initial learning rate of 0.001, a total of 100 epochs, and early stopping implemented after 10 epochs if the validation loss did not improve. The experiments were conducted on two NVIDIA A100 GPUs. The datasets were split into 70% training, 10% validation, and 20% testing sets. We evaluated the performance of three CNN models: ConvNeXt [17], ResNet18 [12], and DenseNet161 [13], using various pooling strategies for image classification. A similar data augmentation procedure to [21, 24, 29] was added to the training data: random cropping and horizontal flipping. In the implementation of the multi-scale lacunarity pooling layer, we use two scales ($S = 2$). The scale was set to this value as the spatial size of the feature maps into the pooling layer were 7×7 .

Table 1. Average test accuracy with ± 1 standard deviation of all pooling method for PlantVillage across three experimental trials. The best average accuracy is bolded for the pooling layers.

Models	Max	Average	L_2	Fractal	Lacunarity Pooling (Ours)		
					Base	DBC	Multi-scale
ConvNeXt	91.71 \pm 0.26	97.30 \pm 0.39	97.30 \pm 0.39	95.74 \pm 0.54	96.48 \pm 0.12	97.71\pm0.17	97.49 \pm 0.30
Resnet18	91.50 \pm 1.26	95.57 \pm 0.21	95.57 \pm 0.17	90.68 \pm 0.50	95.32 \pm 0.44	96.15\pm0.20	95.35 \pm 0.24
Denesenet161	95.47 \pm 0.53	97.49 \pm 0.04	96.32 \pm 0.17	98.01 \pm 0.11	97.66 \pm 0.13	97.75 \pm 0.11	98.07\pm0.21

Table 2. Average test accuracy with ± 1 standard deviation of all pooling method for LeavesTex1200 across three experimental trials. The best average accuracy is bolded for the pooling layers.

Models	Max	Average	L_2	Fractal	Lacunarity Pooling (Ours)		
					Base	DBC	Multi-scale
ConvNeXt	86.88 \pm 1.81	94.88 \pm 1.23	94.88 \pm 1.23	94.44 \pm 0.87	83.22 \pm 2.35	89.22 \pm 1.28	95.00\pm0.82
ResNet18	84.44 \pm 1.09	88.89 \pm 0.31	88.89 \pm 0.31	88.22 \pm 0.16	86.11 \pm 0.87	87.67 \pm 2.32	90.11\pm2.79
DenseNet161	86.33 \pm 1.78	93.44 \pm 1.13	89.78 \pm 1.85	94.11 \pm 1.97	91.11 \pm 0.31	93.22 \pm 2.04	94.55\pm0.31

Table 3. Average test accuracy with ± 1 standard deviation of all pooling method for DeepWeeds across three experimental trials. The best average accuracy is bolded for the pooling layers.

Models	Max	Average	L_2	Fractal	Lacunarity Pooling (Ours)		
					Base	DBC	Multi-scale
ConvNeXt	70.83 \pm 0.71	85.53 \pm 0.78	85.53 \pm 0.78	85.37 \pm 0.29	80.42 \pm 0.57	83.16 \pm 0.20	86.01\pm0.26
ResNet18	66.59 \pm 1.52	73.27 \pm 0.48	73.27 \pm 0.48	63.68 \pm 2.53	69.95 \pm 0.40	69.95 \pm 0.75	73.77\pm0.03
DenseNet161	74.94 \pm 0.75	79.66 \pm 1.09	75.64 \pm 0.42	85.69\pm0.40	74.30 \pm 0.20	75.29 \pm 0.90	80.44 \pm 0.32

4.2. Benchmark Datasets Results

PlantVillage The PlantVillage dataset [14] is for multi-class image classification with 55,448 images divided into 39 classes representing background-only (out of domain images such as animals and buildings), healthy, and diseased plants. Images span 14 plant species: Apple, Blueberry, Cherry, Corn, Grape, Orange, Peach, Bell Pepper, Potato, Raspberry, Soybean, Squash, Strawberry, and Tomato and contains images of 17 fungal diseases, four bacterial diseases, two mold (Oomycete) diseases, two viral diseases, and one disease caused by a mite.

As seen in Table 1, the DenseNet161 model consistently outperformed the other two models, achieving the highest classification accuracy across all pooling methods. Specifically, DenseNet161 with GAP yielded an average accuracy of 97.49%, while our pooling layer yields an improvement in accuracy to 98.07%. Max pooling selects the highest values within the local region. This can be problematic when dealing with images that have varying lighting conditions or shadows such as leaf images since max pooling is sensitive to outliers [24]. In such cases, these extreme values from max pooling might not represent the overall characteristics of the leaf accurately. On the other hand, average pooling or the other pooling layers used calculate “smoothed” values giving a more balanced representation across the local re-

gions and being less influenced by extreme values. The recently proposed fractal pooling technique [10], which aims to capture the fractal nature of textures, has a comparable classification accuracy in our experiments for DenseNet161 and ConvNeXt in comparison to the baseline GAP models, but this comes at a computational cost as discussed in Section 4.6. The DBC method shows an improvement in the accuracy for the ResNet18 and ConvNeXt models without adding additional learnable parameters.

LeavesTex1200 The real-world application, named LeavesTex1200 [4], employs commercial scanner-captured leaf surface images to classify Brazilian plant species. The dataset consists of 20 species, each represented by 60 images, totaling 1200 samples. Specimens were collected in vivo, cleaned, and aligned along the basal/apical axis. From each specimen, three non-overlapping 128×128 pixel windows were extracted to create texture images for classification purposes. As seen in Table 2 the multi-scale lacunarity performs the best when compared to our base lacunarity method. LeavesTex1200 images are distinctly focused on leaves without major variations in texture across a single scale. The leaves have a homogeneous vein structure in the leaves. This lack of textural heterogeneity makes the base lacunarity less effective since the base lacunarity is optimized for capturing textural variations within a single scale. However, the multi-scale lacunarity pooling layer becomes




Figure 5. Example results of DeepWeeds dataset across pooling layers for the Resnet18 model. EigenCAM [22] was used to create the class activation maps to show which areas of the image each model focused on. As we can see in this example, the multi-scale lacunarity pooling layer captures useful information across varying spatial levels of the image.




Figure 6. t-SNE results for LeavesTex1200 Dataset trained on the Densenet161 model. The colors represent the 20 different classes from LeavesTex1200 dataset. Each t-SNE used the same random seed so there is a fair comparison between the different pooling layers. We also computed the average log Fisher Discriminant Ratio (± 1 standard deviation) using the features from the penultimate layer. Higher scores indicate more compact and well separated classes in the higher dimensional space. Multi-scale (MS) lacunarity is shown in Figure 6f. Our quantitative metrics match the qualitative observations in the t-SNE plot.

more adept at capturing subtle texture changes across different scales. When zoomed out, leaves exhibit greater textural heterogeneity, which the multi-scale method effectively leverages, resulting in improved performance compared to both the base lacunarity and baseline average pooling layers.

DeepWeeds DeepWeeds [23] comprises a total of 17,509 RGB images. These images were captured by a specialized ground weed control robot under natural field conditions, without controlled lighting. The dataset covers eight distinct weed species along with various non-weed plants for a total of nine classes. Each weed species is represented by more than 1000 images. DeepWeeds demonstrates better accuracy with the lacunarity pooling layer across ConvNeXt and ResNet18 models as shown in Table 3. However, notable accuracy improvements are particularly evident with the ConvNeXt model. Analysis of the results shows that fractal pooling exhibits better performance, especially with deeper models such as DenseNet161 [10]. Fractal pooling also uses residual block to improve the feature representation captured by this layer. DenseNet161 uses feature concatenation across multiple layers that enables the integration of features representing various levels of abstraction. In texture images, these features correspond to patterns observed at different scales. This multi-scale analysis has long been recognized as crucial for tex-

ture analysis [6], so it is expected that DenseNet models achieves the best accuracy in this context.

Comparing results across all the pooling layers in Tables 1, 2, and 3, it is evident that the lacunarity pooling layers consistently outperform the baseline average pooling layers on average. Specifically, the multi-scale approach demonstrates better performance compared to the base method described in Section 3.2. The DBC method, as discussed in Section 3.3, performs comparably to the average pooling layer. The operations involved in average pooling is similar with L_2 pooling, which is a form of power-average pooling. Given that we focus on feature extraction in our experiments, it is observed that the results for the L_2 pooling closely resemble those obtained through average pooling.

The reason why our lacunarity pooling method, especially the multi-scale variant, outperforms others lies in its capability for capturing the intricate details and varying complexities present in our image data. Unlike typical average pooling techniques that might overlook subtle nuances or textural variations, our multi-scale lacunarity pooling delves into different levels of detail, thereby increasing the model’s ability to discern patterns. Moreover, our pooling approach is finely tuned to use the self-similarity and spatial irregularities inherent in natural textures, which are prevalent in the wide array of plant images we dealt with across different datasets.




Figure 7. Confusion matrices of PlantVillage dataset for fractal (7a) and DBC lacunarity (7b) pooling layers in the ConvNeXt model respectively. Lacunarity improved performance for most classes (34 of 39) demonstrating the effectiveness of our proposed lacunarity pooling layer in comparison to fractal pooling. The average classification accuracy with ± 1 standard deviation is shown below each confusion matrix.

4.3. Explainable AI Analysis

We used an Explainable AI technique, EigenCAM, [22] to visualize the salient regions identified in an input image by the lacunarity pooling layer and other pooling approaches with the ResNet18 model. EigenCAM showed more pronounced areas of significance in the input image, characterized by wider regions of activation for our proposed lacunarity pooling layer fusion model. The higher regions of interest, represented by the brighter highlighted areas, signifies stronger and more focused attention on crucial features within the images. This emphasis on specific regions is advantageous as this aids in capturing intricate details across varying spatial levels of the image. Compared to the other pooling layers such as max, average, and L_2 , we observed improved performance with our model in terms of capturing more information across the input image. This suggests that our proposed lacunarity pooling strategy captures more extensive spatial information, thereby potentially enhancing the discriminative power of the model’s predictions.

4.4. t-SNE Visualization Analysis

The t-Distributed Stochastic Neighbor Embedding (t-SNE) [27] was used to visualize the high-dimensional feature representations learned by our fusion model and other pooling

layers. The same random seed was used to fairly compare each t-SNE projection of the features extracted from each pooling approach. For max pooling as shown in Figure 6b, there is not much separation or compactness in the projected space (similar to L_2 in Figure 6c). However, the lacunarity pooling layers appear to form more compact and separable clusters, especially the multi-scale lacunarity approach.

Moreover, to quantitatively assess the discriminative power of features extracted by our fusion model, we computed the log Fisher Discriminant Ratio (FDR) [24] scores for each class in the original feature space. The FDR score measures the ratio of between-class variation to within-class variation, with higher values indicating greater discriminatory ability. In our experiments, a higher FDR score signifies that the features are more effective at differentiating between classes, thus contributing to enhanced classification performance. When examining the FDR scores across different pooling layers in Figure 6, we observed notable differences. Particularly, the multi-scale lacunarity pooling layer consistently yielded higher FDR scores compared to the DBC lacunarity pooling layer. This suggests that the multi-scale variant of the lacunarity pooling layer captures more discriminative information from the input data, leading to better separability between classes.

4.5. Confusion Matrix Analysis

When comparing the confusion matrices generated by different pooling methods for PlantVillage, we observed improvements in the classification performance of our lacunarity pooling multi-scale approach across a majority of classes. Specifically, out of the total 39 classes in our dataset, our method demonstrated better classification accuracy in 34 classes compared to the fractal pooling layer [10] as seen in Figure 7. For classes like Apple Scab the accuracy increased from 79.84% to 93.65% and Tomato early blight went up from 70.85% to 80.61%. Classes which were misclassified as Corn Cercospora leaf spot instead of Corn Northern leaf blight, where the false positive percentage at 21.50% was reduced to 19.63%. This indicates an enhancement in the model’s ability to correctly classify a wide array of plant diseases and species. The confusion matrices vividly illustrate how our lacunarity pooling method effectively captures the intricate features and subtle variations present in the images, leading to more accurate predictions across diverse classes.

4.6. Model Parameter Analysis

Table 4. Comparison of learnable parameters introduced by fractal and multi-scale (MS) lacunarity pooling layers. All other pooling layers (average, max, L_2 , base lacunarity, and DBC lacunarity) do not introduce additional parameters to the models. Our proposed MS lacunarity pooling layer learns information across multiple levels of the image while being more computationally efficient than the fractal pooling layer.

Model	Fractal	MS Lacunarity
ResNet18	268,297	1,536
ConvNeXt	599,049	2,304
DenseNet161	4,901,769	6,624

Table 4 outlines the number of learnable parameters for each model, highlighting the parameter increase introduced by the fractal method and multi-scale lacunarity across all models. Despite the additional parameters introduced by the fractal method, there is no statistically significant improvement in accuracy compared to the baseline layers except for the DeepWeeds dataset. The multi-scale lacunarity pooling layer introduces extra parameters due to the 1×1 convolution; however, there are improvements in the model’s performance. In the case of DBC for PlantVillage, the DBC lacunarity was a viable alternative to the GAP baseline models to achieve a slight increase in average overall performance without additional learnable parameters.

5. Conclusion

The proposed lacunarity pooling layer is as a novel approach to enhance the performance of CNNs in computer

vision tasks. We presented a new model that integrates lacunarity pooling layers with existing CNN architectures, facilitating the extraction of both local and global spatial information. Through extensive experimentation and analysis, we have demonstrated the efficacy of our approach across multiple datasets and model architectures. Our findings reveal that incorporating lacunarity pooling layers can lead to improved feature extraction and classification accuracy, particularly evident in models such as DenseNet161.

While certain variations in pooling methods gave lower accuracy, our study shows the adaptability and potential of lacunarity-based features in enhancing CNN performance. Future work includes extending this pooling layer to other data modalities (e.g., hyperspectral data), investigating other backbone architectures and adapting the lacunarity pooling layer for other tasks such as segmentation. Another interesting direction is to integrate the multi-scale lacunarity layer with features from different levels of the network. Ultimately, our unique pooling layer demonstrate the potential for the integration of novel textural features like lacunarity as valuable additions to CNN architectures.

6. Acknowledgement

Portions of this research were conducted with the advanced computing resources provided by Texas A&M High Performance Research Computing.

References

- [1] C. Allain and M. Cloitre. Characterizing the lacunarity of random and deterministic fractal sets. *Phys. Rev. A*, 44:3552–3558, 1991. 3
- [2] Francesco Bianconi, Antonio Fernández, Fabrizio Smeraldi, and Giulia Pascoletti. Colour and texture descriptors for visual recognition: A historical overview. *Journal of Imaging*, 7(11), 2021. 1
- [3] Peter J Burt and Edward H Adelson. The laplacian pyramid as a compact image code. In *Readings in computer vision*, pages 671–679. Elsevier, 1987. 4
- [4] Dalcimar Casanova, Jarbas Joaci de Mesquita Sá Junior, and Odemir Martinez Bruno. Plant leaf identification using gabor wavelets. *International Journal of Imaging Systems and Technology*, 19(3):236–243, 2009. 5
- [5] Zhile Chen, Feng Li, Yuhui Quan, Yong Xu, and Hui Ji. Deep texture recognition via exploiting cross-layer statistical self-similarity. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5227–5236, 2021. 2
- [6] CA Coburn and Arthur CB Roberts. A multiscale texture analysis procedure for improved forest stand classification. *International journal of remote sensing*, 25(20):4287–4308, 2004. 6
- [7] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 18(1):1–34, 1999. 1

- [8] P. Dong. Test of a new lacunarity estimation method for image texture analysis. *International Journal of Remote Sensing*, 21(17):3369–3373, 2000. 3
- [9] J. Rex Fiona and J. Anitha. Automated detection of plant diseases and crop analysis in agriculture using image processing techniques: A survey. *2019 IEEE International Conference on Electrical, Computer and Communication Technologies (ICECCT)*, pages 1–5, 2019. 1
- [10] Joao B. Florindo. Fractal pooling: A new strategy for texture recognition using convolutional neural networks. *Expert Systems with Applications*, 243:122978, 2024. 2, 3, 5, 6, 8
- [11] Robert M Haralick, Karthikeyan Shanmugam, and Its’ Hak Dinstein. Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics*, (6):610–621, 1973. 1
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015. 4
- [13] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks, 2018. 4
- [14] David Hughes, Marcel Salathé, et al. An open access repository of images on plant health to enable the development of mobile disease diagnostics. *arXiv preprint arXiv:1511.08060*, 2015. 5
- [15] John E Hutchinson. Fractals and self similarity. *Indiana University Mathematics Journal*, 30(5):713–747, 1981. 1
- [16] James M Keller, Susan Chen, and Richard M Crownover. Texture description and segmentation through fractal geometry. *Computer Vision, Graphics, and Image Processing*, 45(2):150–166, 1989. 2
- [17] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. 4
- [18] Fei Ma, Xuejiao Sun, Fan Zhang, Yongsheng Zhou, and Heng-Chao Li. What catch your attention in sar images: Saliency detection based on soft-superpixel lacunarity cue. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–17, 2023. 2
- [19] Benoit Mandelbrot. How long is the coast of britain? statistical self-similarity and fractional dimension. *Science*, 156(3775):636–638, 1967. 1
- [20] Benoit B Mandelbrot and Benoit B Mandelbrot. *The fractal geometry of nature*. WH freeman New York, 1982. 2
- [21] Akshatha Mohan and Joshua Peeples. Quantitative analysis of primary attribution explainable artificial intelligence methods for remote sensing image classification. In *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2023. 4
- [22] Mohammed Bany Muhammad and Mohammed Yeasin. Eigen-cam: Class activation map using principal components. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7, 2020. 6, 7
- [23] Alex Olsen, Dmitry A. Kononov, Bronson Philippa, Peter Ridd, Jake C. Wood, Jamie Johns, Wesley Banks, Benjamin Girenti, Owen Kenny, James Whinney, Brendan Calvert, Mostafa Rahimi Azghadi, and Ronald D. White. Deep-Weeds: A Multiclass Weed Species Image Dataset for Deep Learning. *Scientific Reports*, 9(2058), 2019. 6
- [24] Joshua Peeples, Weihuang Xu, and Alina Zare. Histogram layers for texture analysis. *IEEE Transactions on Artificial Intelligence*, 3(4):541–552, 2021. 4, 5, 7
- [25] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary Bradski. Kornia: an open source differentiable computer vision library for pytorch. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3674–3683, 2020. 4
- [26] Nirupam Sarkar and Bidyut Baran Chaudhuri. An efficient differential box-counting approach to compute fractal dimension of image. *IEEE Transactions on systems, man, and cybernetics*, 24(1):115–120, 1994. 3
- [27] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. 7
- [28] David P. Williams. Fast unsupervised seafloor characterization in sonar imagery using lacunarity. *IEEE Transactions on Geoscience and Remote Sensing*, 53(11):6022–6034, 2015. 2, 3
- [29] Jia Xue, Hang Zhang, and Kristin Dana. Deep texture manifold for ground terrain recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2018. 4