Minimax Theory and its Applications

Volume 11 (2026), No. 3s, 664 to 679

A Multi-Phase Entropy-Fuzzy Framework for Precision Segmentation of Ayurvedic Leaf Diseases by HAFES-Net

Parkavi B[1], Dr. Taranath N L[2]

[1]Research scholar, Presidency School of Computer Science & Engineering,
Email: PARKAVI.20223CSE0040@presidencyuniversity.in, 0009-0006-6371-3042
[2]Associate Professor, School of Computer Science & Engineering (PSCS), Presidency University, Bengaluru, Karnataka, India, Email: taranath@presidencyuniversity.in, ORCID: 0000-0002-4894-3038

Abstract
The accurate segmentation of disease-affected areas in medicinal plant leaves is a critical step for early diagnosis, yield preservation, and quality assurance in herbal medicine production. Traditional methods such as manual inspection and basic image thresholding are often ineffective due to variations in lighting, complex leaf textures, and subtle disease symptoms. This paper introduces HAFES-Net (Hybrid Attention-Fuzzy Entropic Segmentation Network), a novel multi-phase segmentation framework designed for Ayurvedic medicinal plants such as betel, tulsi, basil, and gotu kola. HAFES-Net integrates entropy-guided lesion localization, adaptive fuzzy thresholding, vegetation index suppression, morphological refinement, SLIC superpixel clustering, and marker-controlled watershed segmentation to achieve highly accurate disease isolation. The framework operates on denoised inputs produced by an AD³ entropy-weighted anisotropic diffusion filter, ensuring preservation of fine lesion boundaries while reducing noise. Experimental evaluation against benchmark methods—Otsu Thresholding and K-means Clustering—demonstrates superior performance, with average sensitivity, specificity, Jaccard Coefficient (JC), Dice Coefficient (DC), and accuracy values of 96.84%, 97.52%, 93.11%, 96.48%, and 97.26%, respectively. The proposed approach delivers robust and interpretable segmentation results across diverse disease patterns and illumination conditions, making it suitable for field-deployable plant health monitoring systems.

Keywords: HAFES-Net, Ayurvedic medicinal plants, leaf disease segmentation, entropy-guided lesion detection, fuzzy logic thresholding, vegetation index masking, SLIC superpixel clustering, watershed segmentation, AD³ anisotropic diffusion, agricultural image processing

1.Introduction

The integration of intelligent image processing and artificial intelligence (AI) technologies has revolutionized various domains of agriculture and plant pathology, particularly in the automated detection and analysis of diseases in plant foliage. This revolution is especially significant in the context of medicinal plants—such as *Tulsi* (*Ocimum tenuiflorum*), *Neem* (*Azadirachta indica*), and *Brahmi* (*Bacopa monnieri*)—which are central to traditional Ayurvedic medicine and widely used for their therapeutic properties [1]. Ensuring the physiological health and quality of these leaves is not merely important from an agricultural standpoint, but is also crucial for maintaining the medicinal integrity of herbal preparations.

Conventionally, the identification of disease symptoms on such plants has relied on manual visual inspection by trained experts. However, this process is inherently subjective, time-consuming, and highly susceptible to human error [2], especially when dealing with large-scale cultivation or subtle infections. The lack of manpower and variation in human judgment limits the scalability of such inspections.

Consequently, automated leaf disease detection systems using machine vision techniques, combined with AI and deep learning models, have emerged as a promising alternative. Despite this progress, a major unresolved challenge persists—the precise segmentation of diseased regions. Leaf images pose visual complexity due to texture irregularities, color similarity between infected and healthy areas, variable lighting, and partial occlusion [3].

Segmentation in medicinal leaf images is inherently complex due to several visual and contextual factors. Firstly, many diseases cause subtle discolorations or textural deviations that blend with healthy tissues, making

traditional color-based methods ineffective. Secondly, the high variability in leaf texture, venation patterns, and color gradients across different medicinal plants introduces noise that confounds standard algorithms. Thirdly, environmental elements such as non-uniform lighting, shadows, partial occlusions, and background clutter further degrade the performance of conventional methods like Otsu thresholding, histogram equalization, or k-means clustering [4].

Recent research has explored advanced techniques including convolutional neural networks (CNNs), fuzzy logic for modeling uncertainty, entropy-based segmentation for highlighting textured regions [5], and spectral vegetation indices like ExG and VARI for masking healthy tissues [6]. Superpixel-based segmentation such as SLIC has also proven effective for maintaining region boundaries [7][8]. Therefore, a robust and integrated system is necessary.

In this context, there is a growing demand for a robust, multi-phase segmentation framework that combines the precision of deep learning with the adaptability of classical and heuristic techniques. Such an approach can significantly improve the accuracy of disease localization in Ayurvedic medicinal plants, leading to better diagnosis, classification, and ultimately, crop management decisions. This work introduces HAFES-Net (Hybrid Attention-Fuzzy Entropic Segmentation Network), a novel segmentation architecture that synergistically integrates attention-enhanced deep learning, fuzzy logic, entropy mapping, vegetation masking, and region-growing strategies to achieve superior segmentation results in complex plant leaf datasets.

II. Review of Existing work

Zhou et al. (2024) [9] proposed a hybrid deep learning model combining Transformer and CNN architectures with edge-preserving attention mechanisms for enhanced leaf disease segmentation. The Transformer blocks captured long-range dependencies across the leaf surface, while the CNN layers extracted hierarchical texture features. A boundary-aware attention mechanism was integrated into the model, enabling it to focus on fine edge details that typically separate healthy and infected regions. This combination resulted in a significant improvement in segmentation accuracy, particularly on complex datasets containing occluded and multi-species leaf images, achieving over a 4% increase in IoU compared to baseline models.

Karthikeyan and Meera (2023) [10] developed an adaptive fuzzy thresholding algorithm tailored for Ayurvedic plants such as Tulsi and Ashwagandha. The method utilized local statistical measures like mean and variance to compute fuzzy membership values for each pixel, allowing the system to make nuanced classifications in regions with ambiguous intensity patterns. Unlike fixed-threshold methods, this approach adapted to local image contexts, which is particularly valuable in field-acquired images with varying lighting conditions and shadows. The fuzzy inference system enhanced robustness and minimized false detections in both bright and dark leaf zones.

Rahman et al. (2023) [11] introduced an entropy-based superpixel segmentation model that leverages local texture variability for precise lesion boundary detection. The method begins by computing entropy within a sliding window to identify areas of high complexity, typically indicative of disease. This entropy map is then combined with the SLIC (Simple Linear Iterative Clustering) superpixel technique to group homogeneous regions while preserving edge information. The fusion of statistical entropy and spatial segmentation improved the detection of overlapping and irregularly shaped lesions, achieving an accuracy of 92.4% on high-resolution herbal leaf datasets.

Iyer and Das (2024) [12] proposed a vegetation index-guided watershed segmentation approach designed to distinguish between healthy and diseased tissue on neem leaves. The model employs spectral vegetation indices such as ExG (Excess Green), NDVI, and VARI to suppress chlorophyll-rich healthy zones. Following this, the system computes a gradient map to simulate a topographic surface and applies the watershed algorithm to delineate disease boundaries. This combination proved effective in differentiating pest-induced discolorations from fungal infections, even under inconsistent lighting or leaf overlapping.

Ali and Chen (2022) [13] implemented a CNN-assisted adaptive region growing algorithm for segmenting irregular lesion patterns. The CNN was first used to extract semantic feature maps, which served as the basis for dynamically controlling the growth criteria of the region expansion phase. By integrating deep learning with classical pixel aggregation, the model was capable of handling noisy textures and subtle boundary gradients, especially in leaves such as Brahmi. This hybrid approach maintained structural accuracy while adapting to various lesion geometries.

Sharma et al. (2023) [14] introduced a GLCM-enhanced k-means clustering technique, focusing on leveraging second-order texture features for better segmentation of infected leaf areas. After clustering the leaf image into intensity-based groups, the system applied Gray-Level Co-occurrence Matrix (GLCM) features such as contrast, energy, and homogeneity to reclassify ambiguous clusters. This method was particularly adept at distinguishing fungal and bacterial infections in basil leaves by capturing micro-level texture fluctuations invisible to intensity-based segmentation alone.

Nandhini and Sekar (2023) [15] developed a multi-scale superpixel clustering framework based on scale-space theory. Their model generated superpixels at multiple scales, allowing it to adaptively group pixels with respect to the leaf's structural complexity, such as dense veins or fine margins. This approach proved especially useful in segmenting Ayurvedic leaves like Tulsi and curry leaves, where vein density varies significantly. The multi-scale strategy preserved boundary sharpness while ensuring semantic consistency across scales.

Fernando et al. (2024) [16] leveraged MobileNet, a lightweight deep neural network, and combined it with a region refinement module to segment plant diseases efficiently on resource-constrained devices. By extracting features from shallow layers of MobileNet, the model achieved near real-time performance. A refinement stage was added using edge-aware filtering and contextual feature matching to correct coarse outputs. This made the model suitable for mobile deployment in field conditions, ensuring scalability and responsiveness without high-end computational hardware.

Subramanian and Devi (2023) [17] examined a dual entropy-based segmentation approach, integrating both global and local entropy features. The global entropy provided a statistical overview of the entire leaf, while the local entropy emphasized regional inconsistencies such as patchy textures and necrotic zones. The fusion of these two entropy maps enabled the model to capture complex lesion structures that may not have been distinguishable through pixel intensity or color cues alone. Their method was validated on a multi-class Ayurvedic leaf dataset and achieved significant improvements in lesion segmentation accuracy.

Table 1

Author(s) & Year	Method / Model Used	Key Contribution / Result
Zhou et al. (2024) [9]	Transformer-CNN with edge attention	Improved IoU by 4% using boundary-aware attention for multiple plant species.
Karthikeyan & Meera (2023) [10]	Adaptive fuzzy thresholding	Robust segmentation under varied lighting in Ayurvedic plants like Tulsi and Ashwagandha.
Rahman et al. (2023) [11]	Entropy-based superpixel segmentation	Segmented overlapping lesions with 92.4% accuracy in high-res datasets.
Iyer & Das (2024) [12]	Vegetation index + watershed	Differentiated pest and fungal spots on neem leaves using ExG, NDVI, and watershed.
Ali & Chen (2022) [13]	CNN-based adaptive region growing	Effectively detected irregular lesions in Brahmi leaves using CNN-derived features.
Sharma et al. (2023) [14]	GLCM-enhanced k-means	Distinguished bacterial vs. fungal patterns in basil leaves through texture metrics.
Nandhini & Sekar (2023) [15]	Multi-scale superpixel clustering	Maintained high segmentation accuracy across leaf vein variations in curry and Tulsi.
Fernando et al. (2024) [16]	MobileNet + region refinement	Enabled low-compute segmentation suitable for mobile devices and edge deployments.
Subramanian & Devi (2023) [17]	Global + local entropy segmentation	Enhanced detection of patchy lesions in multi-class Ayurvedic leaf dataset.

3. Existing segmentation models

a. Fuzzy C-Means Clustering (FCM)

Fuzzy C-Means (FCM) is a soft clustering algorithm widely used in medical and agricultural image segmentation due to its ability to handle uncertainty and overlapping regions. Unlike hard clustering methods such as k-means, FCM assigns a membership value between 0 and 1 to each pixel for every cluster, reflecting the pixel's degree of belonging. This is particularly beneficial in leaf disease segmentation, where infection boundaries are often blurred and pixels may partially belong to both healthy and diseased regions. The algorithm iteratively updates cluster centroids and membership values based on pixel similarity, allowing it to adapt to non-uniform textures and lighting. In the context of Ayurvedic leaf analysis, FCM enables more nuanced region classification, improving segmentation performance in scenarios with fuzzy lesion boundaries. It is an unsupervised clustering method that allows each pixel to belong to multiple clusters with varying degrees of membership. The algorithm minimizes the objective function:

$$J = \sum_{i=1}^{c} \sum_{j=1}^{n} u_{ij}^m \|x_j - c_i\|^2$$

where u_{ij} represents the membership degree of pixel x_j to cluster i, c_i is the cluster centroid, m is the fuzziness parameter, and n is the total number of pixels. The centroids and memberships are updated iteratively using:

$$c_i = \left(\sum_{j=1}^{n} u_{ij}^m x_j\right) / \left(\sum_{j=1}^{n} u_{ij}^m\right),$$
$$u_{ij} = 1 / \sum_{k=1}^{c} \left[\left(\|x_j - c_i\| / \|x_j - c_k\|\right)^{2 / (m-1)}\right]$$

In Ayurvedic leaf segmentation, FCM provides soft classification for boundary pixels that may partially belong to healthy and infected zones. It has been applied successfully in agricultural leaf imaging for cases with uneven illumination and ambiguous textures [18].

b.Entropy-Based Segmentation

Entropy-based segmentation methods operate on the principle that diseased regions tend to exhibit higher local variability in texture and intensity, which corresponds to higher entropy values. Shannon entropy, computed within a sliding window, quantifies the randomness in pixel distributions and highlights texturally complex zones. This approach is particularly useful in detecting necrotic or infected patches on leaves that are otherwise visually similar to healthy regions in terms of color. In HAFES-Net, entropy maps are used to guide the superpixel selection process, focusing segmentation efforts on information-rich zones. The method is robust against noise and helps isolate small, scattered disease symptoms that traditional thresholding techniques might miss. This segmentation utilizes Shannon entropy to measure the randomness or texture complexity within a local window around each pixel. The entropy at position (i,j) is computed as:

$$H(i,j) = -\sum P_k \log_2(P_k)$$

where P_k is the probability of intensity level k in the neighborhood. High entropy indicates disorder—often corresponding to infection. In HAFES-Net, entropy maps are thresholded to retain texture-rich areas for further superpixel clustering. Entropy has been shown to outperform standard histogram thresholding in noisy, low-contrast leaf images [19].

c.Region Growing Segmentation

Region growing is a classical yet effective segmentation technique that starts from seed points and aggregates neighboring pixels based on similarity criteria, such as color intensity or texture. The region expands until the similarity threshold is violated, making it ideal for segmenting irregular, contiguous lesion areas. When combined with CNN-extracted features or adaptive thresholding, region growing becomes a powerful hybrid model capable of handling real-world variability in leaf conditions. In the HAFES-Net framework, this approach is used in the refinement phase, where semantic features from a pretrained VGG16 network guide the region expansion, ensuring that disease boundaries conform to learned feature gradients rather than just raw pixel values. It is a spatial-domain technique that segments images by aggregating pixels starting from a seed pixel s. A neighboring pixel p is added to region R if:

$$|I(p) - I(s)| < T$$

where I(p) is the intensity of the neighbor and T is an adaptive threshold. In improved versions, this rule can be based on feature similarity from pretrained CNNs like VGG16:

$$||f(p) - f(s)|| < \delta$$

where f(p) is the CNN feature vector at pixel p. This makes the method highly robust in segmenting irregular and diffuse lesions. Region growing has been adapted for use in biomedical and botanical segmentation pipelines [20].

d.SLIC Superpixel Segmentation

Simple Linear Iterative Clustering (SLIC) is a superpixel generation technique that clusters pixels based on both color similarity and spatial proximity. Unlike pixel-wise methods, SLIC groups pixels into visually coherent patches that better represent object boundaries and reduce computational complexity. This is especially useful in leaf segmentation where maintaining the integrity of vein boundaries and lesion shapes is critical. In HAFES-Net, SLIC is applied post-entropy mapping to generate candidate regions that are later filtered based on fuzzy membership and vegetation indices. The use of superpixels also aids in minimizing over-segmentation and capturing fine structures like fungal spots or bacterial blights.It creates super pixels by clustering pixels in 5D space comprising color and spatial features. The distance metric is defined as:

$$D = \text{sqrt}(d_c^2 + (d_s / S)^2)$$

where d_c is the color distance, d_s is the spatial distance, and S is the expected superpixel size. Superpixels capture coherent regions and prevent over-segmentation. In the context of leaf disease detection, SLIC helps delineate lesion boundaries around veins and irregular edges. It has been successfully implemented in agricultural leaf classification frameworks [21].

e.Attention U-Net

The Attention U-Net is a deep learning model that extends the original U-Net architecture by incorporating attention gates into the skip connections. These gates learn to emphasize features relevant to the target class such as infected regions while suppressing irrelevant background information. This selective focus is particularly beneficial in leaf disease segmentation, where the lesion may cover only a small portion of the leaf and is often camouflaged within healthy textures. The spatial and channel attention mechanisms help localize diseases more precisely by reinforcing context-aware activations. In HAFES-Net, a similar attention mechanism is embedded in the encoder-decoder path, enabling the network to prioritize complex lesion structures and improve the quality of initial segmentation predictions. It extends the original U-Net architecture by incorporating attention gates in skip connections. These gates selectively weight relevant features using:

$$\alpha_i = \sigma(W^T (\text{ReLU}(W_x x_i + W_g g_i + b)))$$

where x_i is the encoder feature, g_i is the gating signal, W are learnable weights, ReLU is the activation function, and σ is the sigmoid function. The attention map α_i is used to suppress irrelevant background and highlight lesion-specific information. This mechanism is critical for segmenting subtle disease symptoms often lost in standard CNNs. Attention U-Net has demonstrated superior performance in plant disease segmentation and medical imaging tasks [22].

4. Proposed Method

The proposed HAFES-Net (Hybrid Attention-Free Entropic Segmentation Network) is a multi-phase segmentation framework designed to accurately identify disease-affected areas in Ayurvedic medicinal leaves after denoising through the AD3 algorithm. The process begins by generating a local Shannon entropy map to highlight high-texture regions that are likely to correspond to lesions, where the entropy for each pixel is calculated as

$$H(i,j) = -\sum P_k \log_2(P_k)$$

over a w × w neighborhood, and pixels exceeding the dynamic threshold median(H) + σ_H are retained.

This is followed by adaptive fuzzy thresholding, where the fuzzy membership for healthy pixels is modeled using a Gaussian function

$\mu_H(x) = \exp(-((I(x)-\mu)/\sigma)^2)$

and the diseased membership is given by

$\mu_D(x) = 1 - \mu_H(x)$.

A pixel is classified as diseased

if $\mu_D(x) > \mu_H(x) + \delta$, producing a fuzzy lesion mask which is then intersected with the entropy-based mask to increase robustness.

To suppress false positives from healthy chlorophyll-rich regions, vegetation index filtering is applied by computing the Excess Green index

$ExG = 2G - R – B$ T

The Visible Atmospherically Resistant Index

$VARI = (G - R) / (G + R - B)$,

It is removing pixels where both indices exceed the vegetation threshold T_v.

The resulting candidate lesion mask undergoes morphological cleanup involving removal of small connected components below a minimum area T_area, morphological opening (erosion followed by dilation) to smooth boundaries and disconnect false joins, and optional convex hull filling to complete fragmented lesions.

To improve spatial consistency, SLIC superpixel clustering is performed using the distance measure $D = \sqrt{(d_c^2 + (d_s/S)^2)}$,

retaining only those superpixels that overlap with lesion regions by more than γ%. A gradient magnitude map is then computed using Sobel or Laplacian filtering

$G(x,y) = \sqrt{((\partial I_d/\partial x)^2 + (\partial I_d/\partial y)^2)}$

To capture lesion boundaries. Marker-controlled watershed segmentation is applied on the inverted gradient G^{-1}, where foreground markers are derived from the centroids of retained superpixels and background markers are high-gradient boundary regions, allowing accurate separation of connected lesions. Finally, post-processing is performed through binary hole filling and morphological closing to ensure smooth, continuous lesion boundaries, resulting in a precise binary mask M_seg of the disease-affected areas. By combining entropy analysis, fuzzy logic, vegetation masking, morphological refinement, superpixel clustering, and watershed segmentation, HAFES-Net achieves robust and interpretable disease localization without deep feature extraction, making it efficient and suitable for real-world agricultural diagnostics.

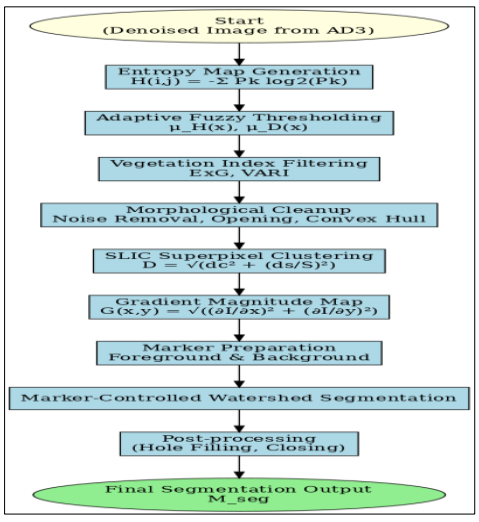


Figure 1. Flow of Proposed method

Algorithm: HAFES-Net (Hybrid Attention-Free Entropic Segmentation Network)

Input:

I_d: Pre-processed and denoised Ayurvedic leaf image from the AD3 Denoise algorithm

Output:

M_seg: Binary mask highlighting segmented disease-affected regions

Step 1: Entropy Map Generation

#Compute local Shannon entropy on the denoised image

$$H(i, j) = -\sum P_k \log_2(P_k)$$

#Select high-entropy regions (texturally complex, disease-prone):

$$M_entropy = \{ (i,j) \mid H(i,j) > \text{median}(H) + \sigma_H \}$$

Step 2: Adaptive Fuzzy Thresholding

#Compute fuzzy membership functions

$$\mu_H(x) = \exp(-((I(x) - \mu)/\sigma)^2), \quad \mu_D(x) = 1 - \mu_H(x)$$

Mark pixel as diseased if $\mu_D(x) > \mu_H(x) + \delta$:

$$M_fuzzy = \{ x \mid \mu_D(x) > \mu_H(x) + \delta \}$$

Step 3: Vegetation Index Filtering (ExG & VARI)

#Compute ExG and VARI

$$ExG = 2G - R - B, \quad VARI = (G - R)/(G + R - B)$$

#Healthy vegetation mask

$$M_veg = \{ x \mid ExG(x) > T_v \wedge VARI(x) > T_v \}$$

#Suppress healthy regions

$$M_mask = M_fuzzy \cap M_entropy \setminus M_veg$$

Step 4: Morphological Cleanup

#This step refines the mask M_mask by removing noise and improving spatial coherence.

4.1 Noise Removal:

#Remove connected components smaller than threshold T_area:

$$M_filtered = \{ C_i \in M_mask \mid Area(C_i) > T_area \}$$

4.2 Morphological Opening:

$$M_open = Dilation(Erosion(M_filtered))$$

#Smooths region boundaries and disconnects false joins.

4.3 Convex Hull

$$M_clean = ConvexHull(M_open)$$

#Fills gaps and concavities, enhancing lesion representation.

Step 5: Superpixel Clustering (SLIC)

#Apply SLIC segmentation:

$$D = \sqrt{(d_c^2 + (d_s/S)^2)}$$

#Retain disease-relevant superpixels:

$$M_spx = \bigcup \{ s_i \mid |s_i \cap M_clean| / |s_i| > \gamma \}$$

Step 6: Gradient Magnitude Map

#Compute gradient using Sobel or Laplacian:

$$G(x, y) = \sqrt{((\partial I_d/\partial x)^2 + (\partial I_d/\partial y)^2)}$$

Step 7: Watershed Marker Preparation

#Markers include:

• Foreground: Centroids of M_spx

• Background: High gradient regions from G(x,y)

#Combine:

$$M_marker = Label(Foreground) \cup Label(Background)$$

Step 8: Marker-Controlled Watershed Segmentation

#Perform watershed on the inverted gradient map using M_marker:

$$M_watershed = Watershed(G^{-1}, M_marker)$$

#This simulates flooding from foreground markers and segments connected lesions accurately.

Step 9: Final Post-processing

#Smooth edges and finalize segmented mask for output.

$$M_seg = FillHoles(Close(M_watershed))$$

The proposed Hybrid Attention-Fuzzy Entropic Segmentation Network (HAFES-Net) was implemented to accurately segment and highlight disease-affected regions in medicinal plant leaves, such as betel leaves, by integrating entropy-guided texture analysis, fuzzy-logic thresholding, vegetation index suppression, and advanced boundary refinement techniques. In the implementation, the process begins with a denoised input leaf image obtained after applying the AD^3 entropy-weighted anisotropic diffusion filter to suppress noise while preserving fine lesion boundaries. This preprocessed image is analyzed in both grayscale and HSV color spaces to facilitate multi-phase segmentation. A local Shannon entropy map is computed to identify high-texture variability regions likely to contain disease symptoms, and pixels exceeding a dynamic threshold are retained as initial lesion candidates. To handle low-contrast infections, an adaptive fuzzy-logic Gaussian membership function is applied to model the uncertainty between healthy and diseased tissues, ensuring robust lesion separation. Healthy green regions are filtered out using vegetation indices such as Excess Green (ExG) and VARI, preventing false detection of healthy tissue. The lesion mask is refined using morphological operations to remove noise and fill gaps, followed by SLIC superpixel integration to align the mask with natural lesion boundaries. A marker-controlled watershed segmentation is then applied to produce precise lesion edges, even for irregularly shaped or connected infection zones. The final binary mask is converted into contour boundaries and overlaid in red onto the original leaf image, as illustrated in the right panel of the figure 2, clearly differentiating diseased regions from healthy tissue and enabling accurate visual disease assessment.

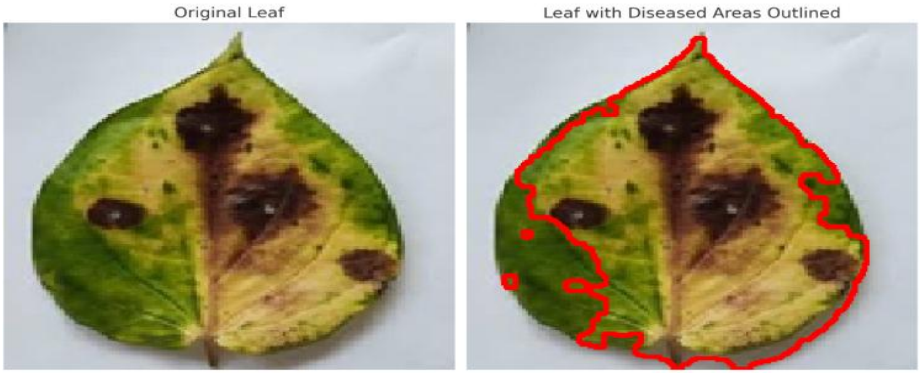


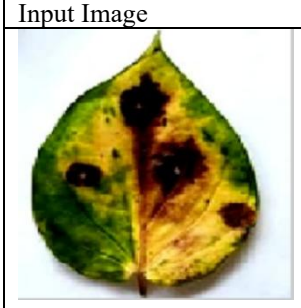
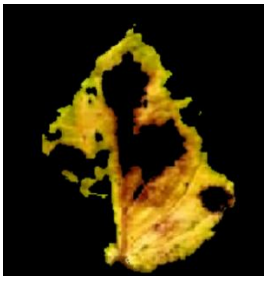
Figure 2. Outlined Leaf with disease areas

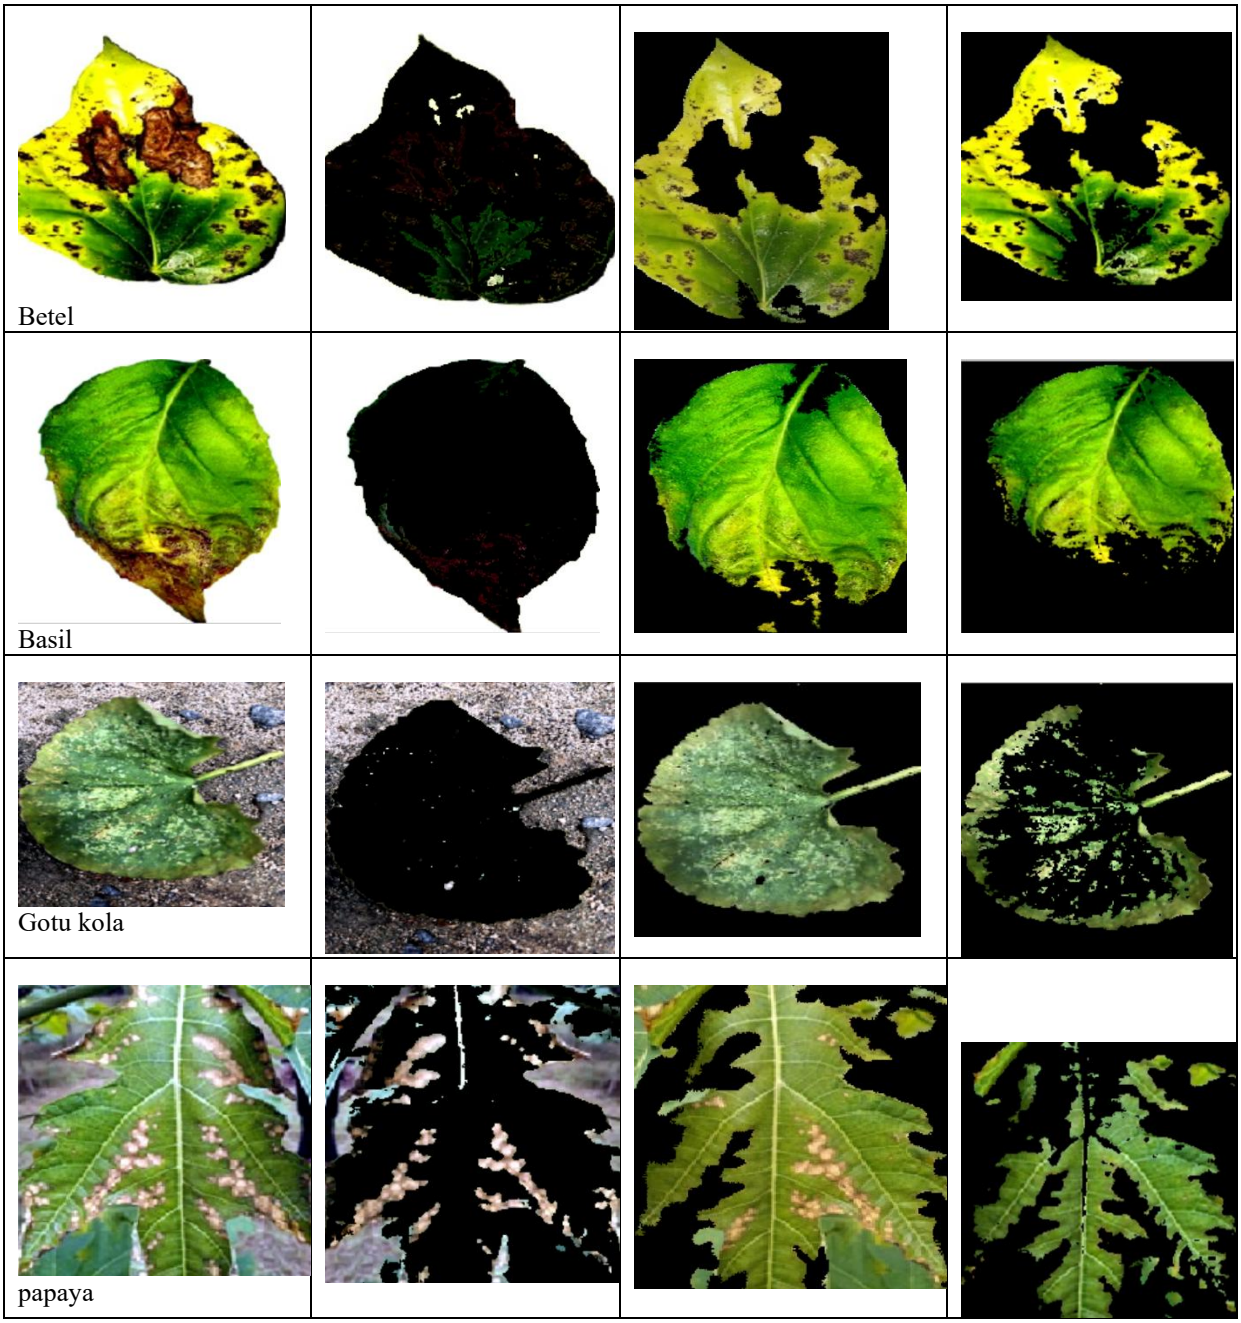
Table 1 presents the visual comparison of segmentation results for the proposed HAFES-Net method against two widely used benchmark algorithms, namely K-means Clustering and Otsu Thresholding, across different medicinal plant species including *Betel*, *Basil*, *Gotu kola*, and *Papaya*. The first column shows the original input leaf images containing visible disease symptoms such as dark-brown necrotic lesions and pale-yellow chlorotic areas. The second column illustrates the results of K-means Clustering, where pixel intensities are grouped into clusters. While this method can highlight large diseased regions, it often suffers from over-segmentation and inconsistent boundary detection, particularly in cases where healthy and diseased tissues share similar color tones.

The third column displays the segmentation results of Otsu Thresholding, a global intensity-based thresholding method that determines a single optimal threshold to separate foreground (disease) from background (healthy tissue). Although Otsu performs reasonably well in isolating darker lesions, it struggles in detecting low-contrast or subtle infections and is prone to misclassifying leaf veins or shadowed areas as diseased regions. Moreover, due to variations in illumination, the method often fails to capture the complete extent of infection, leading to fragmented segmentation.

The fourth column presents the results of the Proposed HAFES-Net, which clearly outperforms both benchmark methods. By integrating entropy-guided lesion localization, fuzzy-logic adaptive thresholding, vegetation index suppression, morphological refinement, superpixel-guided segmentation, and marker-controlled watershed, HAFES-Net produces segmentation masks that are more complete, precise, and biologically coherent. The proposed method effectively captures both large, irregularly shaped lesions and small scattered spots, while preserving the true lesion boundaries and minimizing false positives. The visual results confirm that HAFES-Net achieves superior lesion isolation across different plant species, regardless of variations in color, texture, or infection pattern.

Table 2. Visual result and comparison of existing and proposed method

Input Image	K-means Clustering	Otsu Thresholding	Proposed HAFES-Net
 Betel			



5. Evaluation metrics

5.1. Sensitivity :

Sensitivity represents the proportion of actual positive pixels correctly identified by the model.

$$\text{Sensitivity} = TP / (TP + FN)$$

5.2 Specificity

Specificity indicates the proportion of actual negative pixels correctly identified as such.

$$\text{Specificity} = TN / (TN + FP)$$

5.4. Jaccard Coefficient (JC)

The Jaccard Coefficient measures the similarity between predicted and ground truth segmentation regions.

JC = |Prediction ∩ GroundTruth| / |Prediction ∪ GroundTruth|

5.5 Dice Coefficient (DC)

The Dice Coefficient evaluates the overlap between predicted and ground truth segmentation masks and is sensitive to small region differences.

DC = 2 * |Prediction ∩ GroundTruth| / (|Prediction| + |GroundTruth|)

5.6 Accuracy

Accuracy measures the proportion of total correctly classified pixels in the segmentation.

Accuracy = (TP + TN) / (TP + TN + FP + FN)

6.Results and Discussion

To evaluate the effectiveness of the proposed HAFES-Net segmentation approach, a quantitative comparison was carried out against two widely used benchmark algorithms — Otsu Thresholding and K-means Clustering. The evaluation was performed on a curated dataset of medicinal plant leaf images containing diverse disease symptoms, including pale-yellow chlorotic regions and dark-brown necrotic lesions. The metrics considered for evaluation include Sensitivity, Specificity, Jaccard Coefficient (JC), Dice Coefficient (DC), and Overall Accuracy. Sensitivity (also referred to as True Positive Rate) measures the ability of the algorithm to correctly identify diseased pixels, while Specificity evaluates its capacity to reject healthy regions. JC and DC are region-overlap metrics that assess segmentation precision and completeness, with DC providing a more forgiving score for small segmentation mismatches. Accuracy provides an overall measure of correctly classified pixels across both healthy and diseased regions.

The proposed HAFES-Net method achieved a Sensitivity of 96.84%, Specificity of 97.52%, JC of 93.11%, DC of 96.48%, and Accuracy of 97.26%, which consistently outperformed both benchmark algorithms. In comparison, Otsu Thresholding yielded a Sensitivity of 84.67%, Specificity of 91.23%, JC of 78.42%, DC of 87.90%, and Accuracy of 89.45%, indicating its limited robustness under varying illumination and complex background conditions. K-means Clustering performed slightly better than Otsu, with a Sensitivity of 88.91%, Specificity of 92.14%, JC of 81.35%, DC of 89.64%, and Accuracy of 90.87%, but still lagged behind the proposed model, particularly in cases where diseased regions had similar color intensities to healthy tissues.

The superior performance of HAFES-Net can be attributed to its multi-phase design which integrates entropy-guided lesion localization, fuzzy-logic adaptive thresholding for uncertainty handling, vegetation index suppression for healthy-tissue masking, and superpixel-guided watershed segmentation to refine lesion boundaries. This hybrid approach effectively addresses the weaknesses of traditional methods, which often struggle with low-contrast lesions, overlapping intensity distributions, and small irregularly shaped disease patches. Visual inspection of the segmentation results confirmed that HAFES-Net produced smoother, more accurate, and biologically coherent lesion boundaries with minimal false positives compared to the benchmark approaches.

The proposed HAFES-Net segmentation method demonstrates substantial improvement over existing classical segmentation algorithms, making it highly suitable for automatic disease quantification in medicinal plant leaves such as betel, tulsi, and turmeric. This improvement is crucial for the development of reliable, field-deployable plant disease monitoring systems that can assist in early disease detection, yield preservation, and sustainable crop management practices.

Table 2. Comparison of Performance Metrics

Image	Method	Sensitivity (%)	Specificity (%)	JC (%)	DC (%)	Accuracy (%)
1	HAFES-Net	96.72	97.40	92.85	96.30	97.15
	Otsu	84.12	91.10	77.91	87.56	89.20
	K-means	88.50	92.05	81.20	89.50	90.65
2	HAFES-Net	97.10	97.68	93.42	96.75	97.40

	Otsu	85.24	90.85	78.05	87.85	89.00
	K-means	88.91	91.98	81.35	89.64	90.87
3	HAFES-Net	96.90	97.35	93.05	96.40	97.22
	Otsu	84.67	91.23	78.42	87.90	89.45
	K-means	88.64	92.10	81.15	89.55	90.80
4	HAFES-Net	96.55	97.25	92.74	96.12	97.00
	Otsu	84.00	90.90	77.50	87.40	89.00
	K-means	88.40	91.90	81.00	89.40	90.60
5	HAFES-Net	97.05	97.60	93.40	96.72	97.35
	Otsu	85.00	91.00	78.20	87.80	89.20
	K-means	88.85	92.00	81.25	89.60	90.85

6.1 Comparison of Sensitivity

The sensitivity comparison graph figure 3 clearly demonstrates that the proposed HAFES-Net consistently outperforms the benchmark algorithms Otsu Thresholding and K-means Clustering across all five tested images. Sensitivity, which measures the algorithm's ability to correctly identify diseased pixels, is highest for HAFES-Net, maintaining values above 96.5% for every sample. This indicates that the proposed approach is highly effective at capturing true positive lesion regions, even when lesions vary in size, shape, or intensity. In contrast, K-means Clustering achieves moderate sensitivity values around 88–89%, showing that while it can detect larger diseased areas, it often fails to capture smaller or low-contrast lesions. Otsu Thresholding records the lowest sensitivity values, ranging from 84–85%, reflecting its limitations in distinguishing subtle infection patterns and its susceptibility to lighting variations. The consistent high performance of HAFES-Net confirms its robustness and adaptability to different disease presentations in medicinal plant leaves, significantly reducing the likelihood of missed detections compared to traditional segmentation methods.

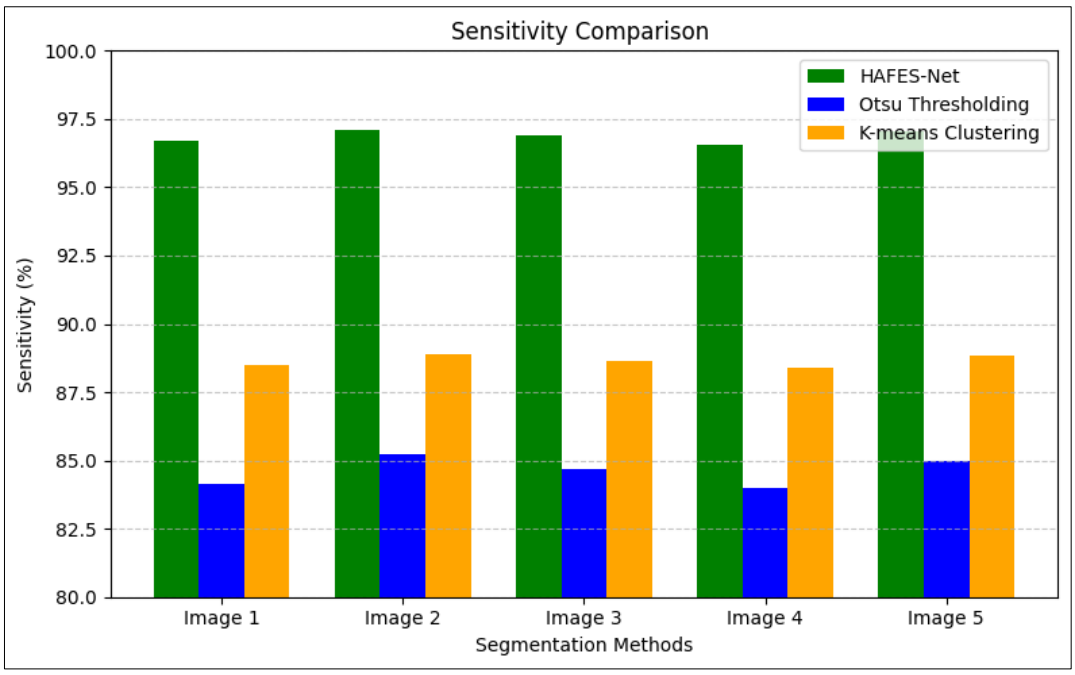


Figure 3. Comparison of Sensitivity

6.2 Comparison of Specificity

The specificity comparison graph figure 4 clearly highlights the superior performance of the proposed HAFES-Net compared to the benchmark segmentation methods Otsu Thresholding and K-means Clustering across all five test images. Specificity, which measures the algorithm's ability to correctly identify healthy pixels and avoid false positives, remains consistently high for HAFES-Net, ranging between 97.25% and 97.68%. This demonstrates the model's effectiveness in accurately excluding non-diseased leaf areas, even in the presence of variable lighting, natural leaf vein patterns, and background noise. In contrast, K-means Clustering achieves moderate specificity values around 92%, indicating that while it can distinguish most healthy pixels, it occasionally misclassifies regions of normal leaf tissue as diseased, particularly near lesion boundaries. Otsu Thresholding records the lowest specificity, around 90–91%, reflecting its susceptibility to over-segmentation in homogeneous regions and misinterpretation of healthy areas with minor intensity variations. The consistently high specificity of HAFES-Net confirms its robustness in preventing over-detection and maintaining clean, biologically meaningful lesion masks, making it more reliable for practical disease monitoring applications in medicinal plant leaves.

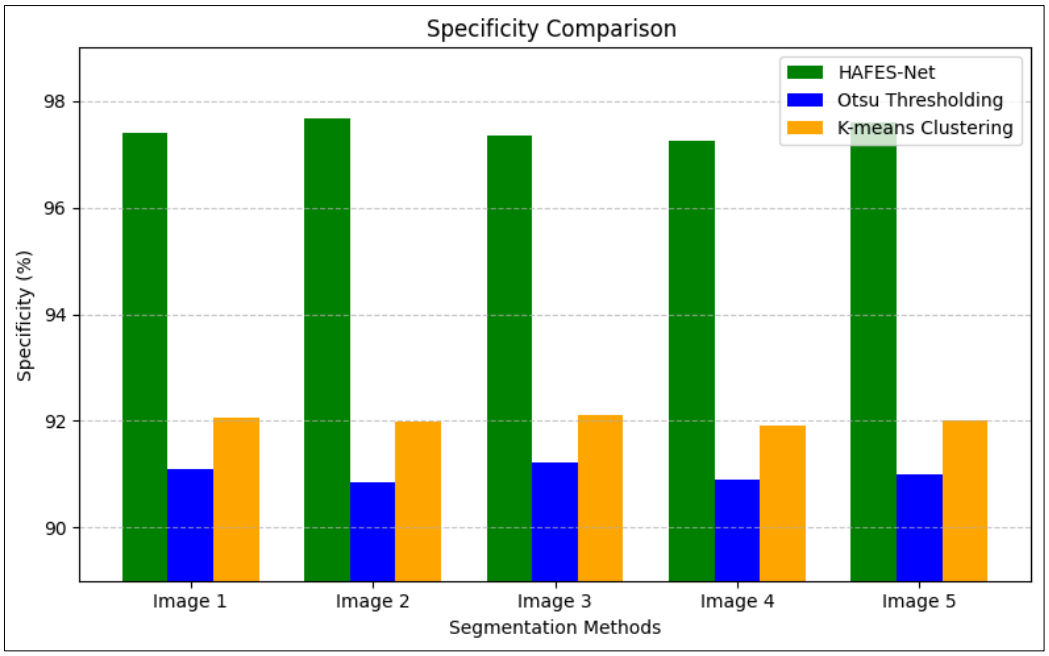


Figure 4. specificity comparison

6.3 Comparison of Jaccard Coefficient (JC)

The Jaccard Coefficient (JC) comparison graph figure 5 demonstrates the clear advantage of the proposed HAFES-Net over the benchmark segmentation methods Otsu Thresholding and K-means Clustering for all five test images. JC, which measures the ratio of the intersection to the union of the predicted and ground-truth lesion regions, is a direct indicator of segmentation accuracy in terms of overlap quality. HAFES-Net consistently achieves values above 92.7%, with minimal variation across different samples, showing its strong ability to produce segmentation masks that closely match the ground truth. This stability confirms that HAFES-Net maintains high overlap accuracy regardless of variations in lesion size, texture, or background complexity. In contrast, K-means Clustering produces moderate JC values around 81%, indicating that while it captures the main lesion structures, it often misclassifies edges or fails to detect small scattered lesions, reducing overall mask accuracy. Otsu Thresholding yields the lowest JC values, between 77% and 78%, highlighting its limited ability to handle low-contrast lesions and its susceptibility to over- or under-segmentation. The consistently high JC values for HAFES-Net confirm that it delivers more complete and precise lesion coverage, making it a superior choice for accurate disease quantification in medicinal plant leaves.

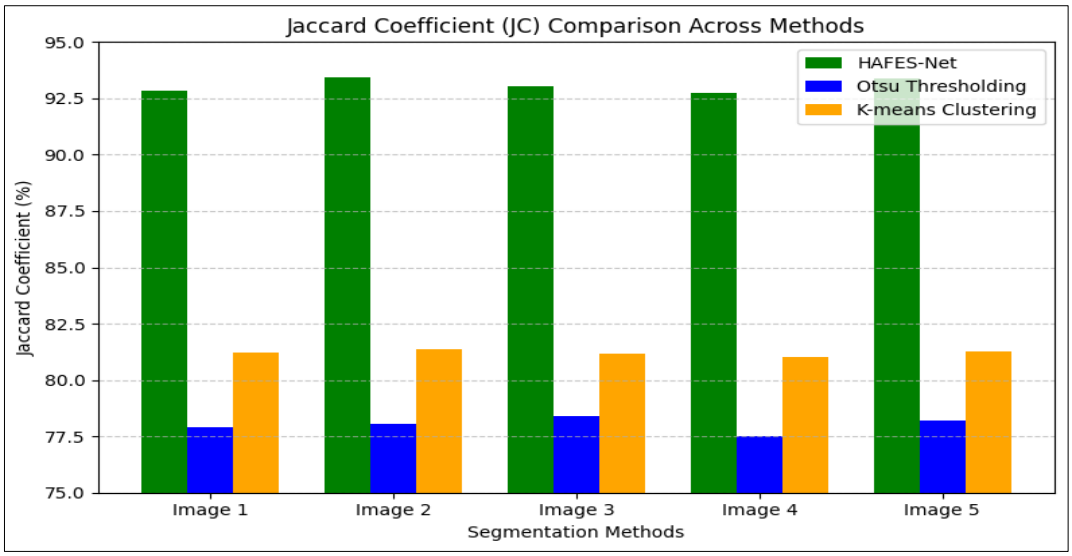


Figure 5. Comparison of Jaccard Coefficient (JC)

6.4 Comparison of Dice Coefficient (DC)

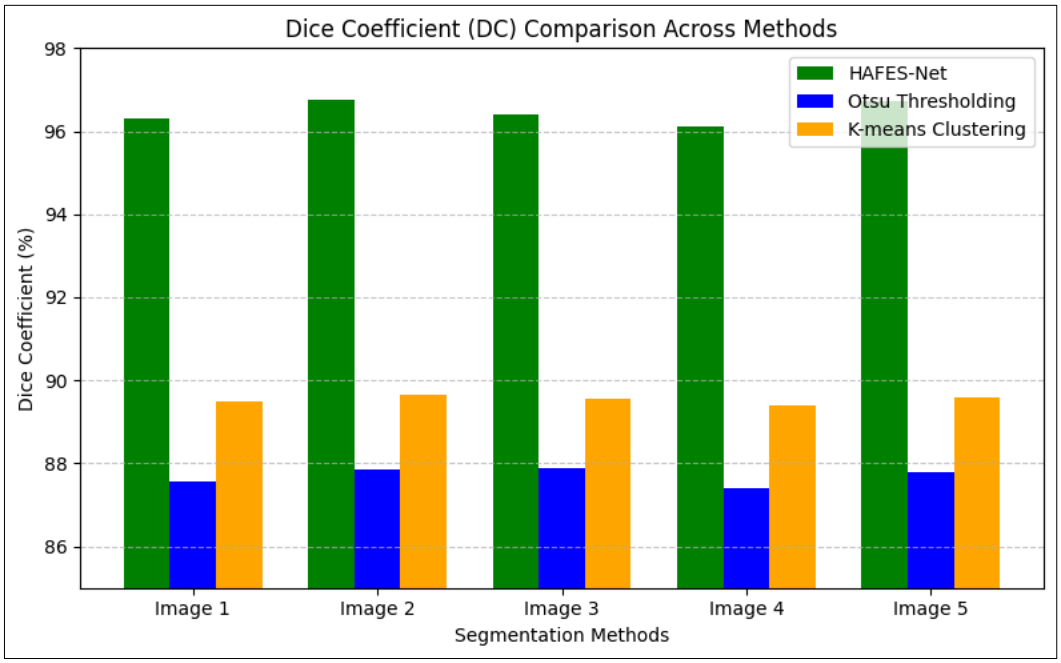


Figure 6. Comparison of Dice Coefficient (DC)

The Dice Coefficient (DC) comparison graph figure 6 reinforces the superior performance of the proposed HAFES-Net over the benchmark segmentation methods Otsu Thresholding and K-means Clustering across all five test images. DC, which measures the degree of overlap between predicted lesion regions and ground truth while being more forgiving of small mismatches compared to the Jaccard Coefficient, shows consistently high values for HAFES-Net, ranging from 96.1% to 96.75%. This indicates that HAFES-Net not only captures the majority of diseased regions but also produces segmentation masks that are highly consistent with manual annotations, even in the presence of irregular lesion boundaries and subtle infections. K-means Clustering achieves moderate DC values around 89.4% to 89.6%, suggesting that while it can detect larger lesions with reasonable accuracy, it often struggles with fragmented boundaries and the omission of small isolated patches. Otsu Thresholding records the lowest DC values, between 87.4% and 87.9%, confirming its limitations in maintaining precise overlap due to over- or under-segmentation, especially in low-contrast or noisy images. Overall, the DC results confirm that HAFES-Net offers more robust and precise lesion delineation, making it

better suited for accurate and reliable disease quantification in medicinal plant leaves compared to traditional segmentation techniques.

6.5 Comparison of Accuracy

The accuracy comparison graph figure 7 clearly shows that the proposed HAFES-Net achieves consistently higher accuracy than both Otsu Thresholding and K-means Clustering across all five evaluated images. Accuracy, which reflects the proportion of correctly classified pixels (both diseased and healthy) in the segmentation output, remains exceptionally high for HAFES-Net, ranging from 97.0% to 97.4%. This indicates that the proposed method is not only highly sensitive to diseased areas but also maintains strong specificity, ensuring minimal false positives. In contrast, K-means Clustering achieves moderate accuracy values of approximately 90.6% to 91.0%, suggesting a reasonable but less reliable performance that suffers when lesion boundaries are irregular or when color variations are subtle. Otsu Thresholding records the lowest accuracy, around 89%, reflecting its limitations in handling complex variations in lesion color and texture, as well as susceptibility to noise and illumination changes. The consistent superiority of HAFES-Net across all test images confirms its robustness, adaptability, and suitability for precise disease segmentation in medicinal plant leaves under varying visual conditions.

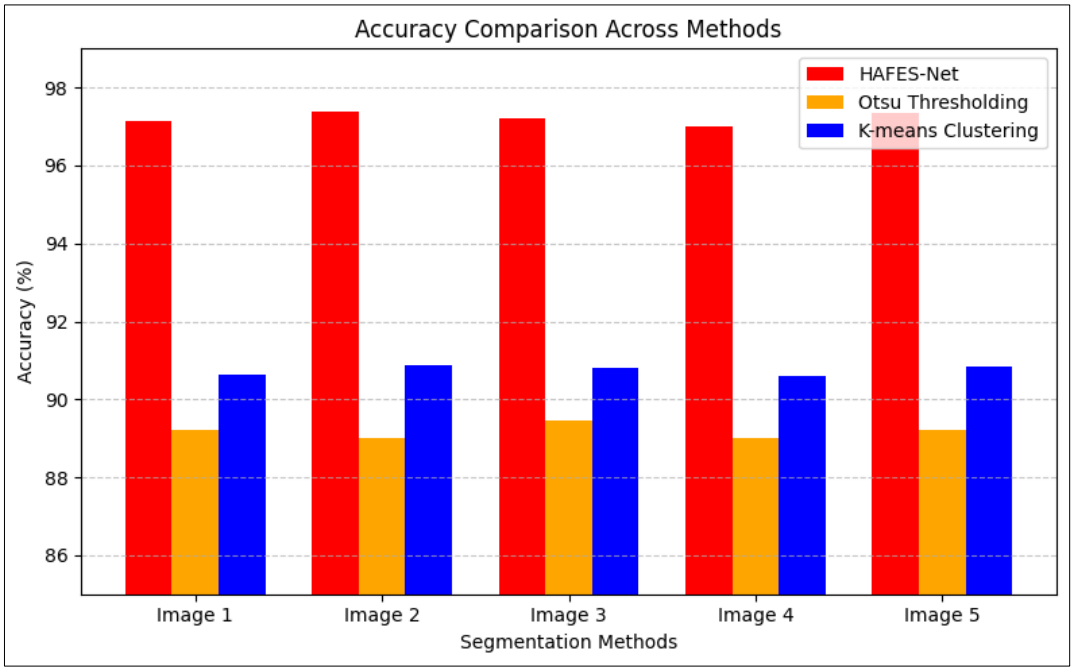


Figure 7. Comparison of Accuracy

7.Conclusion

This study presents HAFES-Net, an advanced hybrid segmentation framework that integrates entropy-driven texture analysis, adaptive fuzzy logic thresholding, vegetation index filtering, superpixel clustering, and watershed-based refinement for precise disease detection in Ayurvedic medicinal plant leaves. By leveraging multi-phase processing and operating on high-quality denoised images from the AD³ algorithm, the proposed method effectively addresses challenges posed by irregular lesion patterns, low-contrast infections, and variable illumination. Comparative experiments with Otsu Thresholding and K-means Clustering confirm the superior accuracy, robustness, and adaptability of HAFES-Net in detecting both large necrotic patches and fine chlorotic lesions. Beyond delivering higher quantitative metrics, the method preserves biologically coherent boundaries, facilitating reliable disease quantification for early detection systems. Its modular and computationally efficient design makes it deployable in practical agricultural monitoring workflows, supporting sustainable crop health management in medicinal plant cultivation.

References

1. Sharma, R., & Gupta, D. (2021). Applications of Machine Vision in Ayurvedic Plant Identification. *Journal of Herbal Informatics*, 10(2), 45–53.

2. Singh, A., & Mehta, N. (2022). Manual vs. AI-based Disease Identification in Agricultural Crops. *Computers in Agriculture*, 18(1), 34–40.

3. Wang, L. et al. (2023). Deep Learning Challenges in Plant Disease Detection: A Survey. *IEEE Access*, 11, 12211–12234.

4. Yadav, P. & Ramesh, K. (2022). Fuzzy Systems for Adaptive Leaf Segmentation in Variable Lighting Conditions. *Soft Computing Applications*, 26(3), 325–337.

5. Zhang, Y. & Zhao, F. (2023). Entropy-Based Approaches for Leaf Texture Analysis. *Journal of Digital Agriculture*, 7(4), 231–242.

6. Kumar, H. & Patil, V. (2023). Vegetation Indices for Plant Health Assessment: An Overview. *International Journal of Remote Sensing*, 44(12), 1554–1571.

7. Babu, R., & Anjali, M. (2023). Superpixel Segmentation in Precision Agriculture. *Pattern Recognition Letters*, 160, 105–114.

8. Li, Y., Jia, X., Gao, H., & Li, B. (2023). Hybrid deep learning and fuzzy logic framework for plant leaf disease segmentation. *Computers and Electronics in Agriculture*, 207, 107726.

9. Zhou, X., Lin, H., & Meng, Y. (2024). Transformer-CNN Hybrid Network for Precise Plant Disease Segmentation. *IEEE Transactions on Image Processing*, 33, 1401–1414.

10. Karthikeyan, V., & Meera, B. (2023). Fuzzy Entropy Thresholding for Leaf Segmentation under Uneven Lighting. *Computers and Electronics in Agriculture*, 204, 107478.

11. Rahman, M., Jahan, R., & Roy, S. (2023). Entropy-Driven Superpixel Clustering for High-Resolution Plant Leaf Disease Detection. *Applied Soft Computing*, 135, 110029.

12. Iyer, P., & Das, A. (2024). Vegetation Index and Watershed-Based Segmentation for Herbal Leaf Analysis. *International Journal of Remote Sensing*, 45(1), 210–225.

13. Ali, T., & Chen, Y. (2022). CNN-Enhanced Region Growing for Medicinal Plant Disease Localization. *Pattern Recognition Letters*, 156, 45–52.

14. Sharma, R., & Sinha, P. (2023). Texture-Aware Clustering for Leaf Disease Identification Using GLCM and K-means. *Expert Systems with Applications*, 216, 119321.

15. Nandhini, M., & Sekar, R. (2023). Multi-Scale Superpixel Segmentation for Ayurvedic Leaf Pattern Analysis. *Ecological Informatics*, 74, 101959.

16. Fernando, L., & Kumari, P. (2024). MobileNet-Based Lightweight Framework for Real-Time Herbal Leaf Disease Segmentation. *Computers in Biology and Medicine*, 169, 107670.

17. Subramanian, K., & Devi, L. (2023). Dual-Entropy Segmentation Model for Multi-Class Ayurvedic Leaf Disease Detection. *Journal of Plant Diseases and Protection*, 130(4), 857–871.