



Article

Decoupled Geometric Measurement and Machine Learning Classification for Automated Post-Harvest Quality Assessment of *Ruscus hypophyllum* Foliage

Fernando Ortega-Loza ^{1,*} , Fernando Toapanta-Ramos ² , Diego Peña ^{3,4} , Moad Hicham Safhi ⁵ and Diego H. Peluffo-Ordóñez ⁶

- ¹ Departamento de Ciencias de la Energía y Mecánica, Universidad de las Fuerzas Armadas ESPE, Av. Gral. Rumiñahui S/N, Sangolquí 171103, Ecuador
- ² Mechanical Engineering, Research Group in Engineering, Productivity and Industrial Simulation (GIIPSI), Master's Program in Mechanical Engineering, Universidad Politécnica Salesiana, Quito 170702, Ecuador; ltoapanta@ups.edu.ec
- ³ Faculty of Engineering Sciences, Universidad Técnica Estatal de Quevedo, Av. Quito km. 1 1/2 vía a Santo Domingo de los Tsáchilas, Quevedo 120301, Ecuador; dpena@uteq.edu.ec or dppb0002@red.ujaen.es
- ⁴ Departamento de Ingeniería Eléctrica, Universidad de Jaén, 23700 Jaén, Spain
- ⁵ C3S Laboratory, Higher School of Technology, Hassan II University, Casablanca 20410, Morocco; moadhicham.safhi@univh2c.ma
- ⁶ School of Mathematical and Computational Sciences, Yachay Tech University, Urcuquí 100150, Ecuador; diego.peluffo@sdas-group.com
- * Correspondence: fwortega@espe.edu.ec

Highlights

What are the main findings?

- A hybrid machine vision framework explicitly decouples geometric stem measurements from visual foliage quality classification for *Ruscus hypophyllum*.
- The pixel-based geometric estimation stage achieves a Mean Absolute Error (MAE) of 1.2 mm and a 99.84% size categorization accuracy.
- Out of 24 classifier configurations evaluated, the Linear Support Vector Machine (SVM) achieves the best performance with an accuracy of 92.4%.

What are the implications of the main findings?

- Provides an interpretable, computationally efficient, and non-destructive inspection solution suitable for automated foliage grading systems.
- Offers an accessible technology to optimize quality control and sorting efficiency in export-oriented ornamental horticulture industries.

Abstract

This paper presents a hybrid computer vision framework that explicitly decouples geometric stem measurement from visual foliage condition classification for automated post-harvest quality assessment of *Ruscus hypophyllum* ornamental foliage. The proposed approach addresses a gap in the literature where heterogeneous quality attributes are typically treated within a single unified learning framework. In the first stage, stem size is estimated using a pixel-based geometric method that incorporates trigonometric orientation correction and spatial calibration via a reference marker of known length, enabling accurate conversion of image measurements to real-world physical dimensions. In the second stage, foliage condition is classified as good or poor using supervised machine learning models trained on Bag of Features representations extracted with the SIFT descriptor. A dataset of 1233 *Ruscus hypophyllum* images was acquired under controlled conditions using a consumer-grade smartphone camera and processed using open-source Python 3.11



Academic Editors: Xiu Jin, Long Zhao and Danyan Chen

Received: 18 June 2026

Revised: 9 July 2026

Accepted: 16 July 2026

Published: 29 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

libraries. Twenty-four classifier configurations across six model families were evaluated using stratified 10-fold cross-validation. The geometric estimation stage achieved a Mean Absolute Error (MAE) of 1.2 mm, a Root Mean Square Error (RMSE) of approximately 1.3 mm, and a size categorization accuracy of 99.84%. For foliage condition classification, the Linear Support Vector Machine achieved the best performance, with an accuracy of $92.4 \pm 1.0\%$ and an F1-score of $91.1 \pm 1.2\%$, outperforming all other evaluated configurations. The proposed framework provides an interpretable, computationally efficient, and accessible solution for automated foliage quality grading, with potential applications in export-oriented ornamental foliage processing facilities.

Keywords: post-harvest quality assessment; ornamental foliage; *Ruscus hypophyllum*; machine vision; intelligent inspection; geometric measurement; SIFT descriptor; Support Vector Machine

1. Introduction

Agricultural production and the export of ornamental plants represent an important economic sector in several countries, including Ecuador. The global ornamental horticulture industry—encompassing cut flowers, potted plants, and cut foliage—currently generates more than USD 130 billion annually, involving more than 90 countries worldwide [1]. Within this context, Ecuador has consolidated its position as one of the leading exporters of ornamental plants and cut flowers at the international level, with floriculture constituting one of the most dynamic areas of its national economy in terms of export value and rural employment generation [2]. Among commercially relevant ornamental species, *Ruscus hypophyllum* has gained increasing importance in international floral markets due to its durability, aesthetic characteristics, and suitability for floral arrangements [3]. The ornamental horticulture sector is highly competitive, driven by continuous advances in infrastructure availability, variety development, and post-harvest handling technologies [4]. In post-harvest processing chains, product quality cannot be improved after harvest but only preserved, making consistent and reliable quality assessment a critical factor for competitiveness in global markets [1].

The post-harvest handling of ornamental foliage typically involves several stages, including harvesting, selection, classification, and packaging. In many production systems, these processes are still carried out manually by trained workers who evaluate visual characteristics such as stem length, leaf distribution, and overall plant condition. However, manual inspection presents well-documented limitations in agricultural quality control contexts: it is inherently subjective, prone to inter-operator variability, and inefficient for large-scale production environments [5,6]. In particular, visual assessments performed by different operators on the same batch of products have been shown to yield inconsistent results, compromising the reliability of quality grading [7]. Furthermore, the physical demands of repetitive manual inspection reduce throughput and increase the risk of error [6]. These limitations are especially critical in export-oriented post-harvest chains, where strict and consistent quality standards must be met to satisfy international market requirements [3]. These limitations motivate the development of automated vision-based inspection systems capable of providing objective, repeatable, and scalable quality assessments in post-harvest processing environments.

Recent advances in computer vision and machine learning have enabled the development of automated inspection systems across a wide range of agricultural applications, including fruit grading, plant disease detection, and crop monitoring [6,8–10]. These sys-

tems typically rely on image-based feature extraction combined with supervised learning models to classify agricultural products according to predefined quality criteria [5,7].

Within the specific domain of ornamental plant inspection, computer vision techniques have been applied to the classification and grading of cut flowers. A hybrid vision-based approach for the identification of *Anthurium* flower cultivars was proposed in [11], combining geometric feature extraction with machine learning classifiers and achieving classification accuracies above 99%. More recently, image-based deep learning methods have been explored for quality grading of cut flowers based on maturity status, demonstrating the feasibility of automated ornamental grading under controlled acquisition conditions [12]. Automated trait estimation in ornamental floriculture has also been investigated using convolutional neural networks, yielding accuracies between 35% and 99% depending on the trait evaluated [13]. However, these studies focus primarily on flower-type ornamentals and do not address the specific challenges associated with cut foliage species such as *Ruscus hypophyllum*, where both stem morphology and foliar condition must be jointly assessed.

Regarding the geometric measurement of plant traits, image-based methods have been successfully applied to non-destructive estimation of morphological parameters such as stem length, plant height, and leaf area [14,15]. These approaches typically rely on spatial calibration using reference markers and scale factors to convert pixel measurements into real-world dimensions, as reviewed in the context of precision agriculture and plant phenotyping [16,17]. Despite their demonstrated effectiveness, such geometric methods have not been systematically combined with visual quality classification for post-harvest foliage assessment, representing a gap that the present work seeks to address.

Table 1 summarizes representative studies on vision-based quality assessment of ornamental plants, highlighting the species, attributes evaluated, methodological approach, and reported accuracy. As shown, existing work focuses primarily on flower-type ornamentals and does not simultaneously address geometric and visual quality attributes in cut foliage species such as *Ruscus hypophyllum*.

Table 1. Representative studies on vision-based quality assessment of ornamental plants and related species.

Study	Species	Attributes	Method	Accuracy
Soleimanipour & Chegini [11]	<i>Anthurium</i>	Cultivar identification	Geometric + ML	>99%
Li et al. [12]	Cut flowers	Maturity/quality grading	CNN (4D deep learning)	N/R
Afonso et al. [13]	Ornamental floriculture	Multiple morphological traits	CNN	35–99%
Franco et al. [18]	Pastures	Biophysical parameters	Hybrid ML + CV	85–90%
This work	<i>Ruscus hypophyllum</i>	Stem size and foliar condition	Decoupled geometric + BoF/SVM	99.84% + 92.4%

Classical machine learning models—including Support Vector Machines, Decision Trees, Naive Bayes, and k-Nearest Neighbors—remain competitive and widely used in image classification tasks that involve structured feature representations under limited data conditions [19,20]. In particular, the Bag of Features framework combined with local descriptors such as SIFT has demonstrated robust performance in plant-related classification tasks where deep learning approaches may be constrained by dataset size [21,22]. These characteristics make classical models particularly suitable for the controlled acquisition scenario addressed in this work.

Beyond data volume considerations, the interpretability, algorithmic stability, and lightweight deployment profile of classical machine learning pipelines constitute independent motivations for their adoption in resource-constrained agricultural inspection

environments, where fault traceability, reproducibility of grading outcomes, and low maintenance overhead are critical operational requirements [3,6].

Despite these advances, most existing approaches treat quality assessment as a unified classification task, in which all attributes of interest are learned simultaneously by the model [8,11]. In the case of ornamental foliage such as *Ruscus hypophyllum*, however, quality assessment involves two fundamentally different types of attributes: measurable physical properties—such as stem length—that can be precisely estimated using geometric image analysis, and visual attributes—such as leaf density, color uniformity, and foliar damage—that require learning-based classification. Treating these heterogeneous attributes within a single unified framework introduces unnecessary complexity in three specific dimensions: it increases the amount of labeled training data required to jointly learn geometric and visual patterns, reduces model interpretability by conflating deterministic and statistical processes within a single learned representation, and obscures the availability of deterministic geometric methods that are both more accurate and more transparent than learned models for measurable physical traits [15,18].

To address this gap, this paper proposes a two-stage hybrid computer vision framework for the automated quality assessment of *Ruscus hypophyllum* foliage. In the first stage, stem size is estimated using a geometric pixel-based method that incorporates trigonometric orientation correction and spatial calibration via a reference marker, enabling accurate conversion of image measurements to real-world dimensions [16,17]. In the second stage, foliage condition is classified as good or poor using supervised machine learning models trained on Bag of Features representations extracted with the SIFT descriptor [23,24]. By decoupling these two tasks, the proposed system achieves both high measurement accuracy and interpretable quality classification without requiring a unified end-to-end learning model.

The main contributions of this work are as follows:

- A hybrid two-stage framework that explicitly decouples geometric stem measurement from visual foliage condition classification, improving system interpretability and reducing model complexity compared to unified classification approaches.
- A pixel-based stem size estimation method incorporating trigonometric orientation correction and reference-marker spatial calibration, enabling accurate conversion of pixel measurements to real-world physical dimensions under controlled acquisition conditions.
- A systematic comparative evaluation of classical supervised machine learning classifiers on Bag of Features representations, providing evidence-based guidance on model selection for structured visual feature spaces in agricultural inspection tasks.
- An experimental validation on a curated dataset of 1233 *Ruscus hypophyllum* images acquired under controlled post-harvest conditions, demonstrating the practical applicability of the proposed system in real agricultural processing environments.

The remainder of this paper is organized as follows. Section 2 describes the proposed methodology and the image acquisition process. Section 3 presents the experimental results for both the geometric estimation and classification stages. Section 4 discusses the obtained results, their implications, and the limitations of the proposed approach. Finally, Section 5 presents the conclusions and future research directions.

2. Materials and Methods

The proposed system follows a hybrid approach that combines geometric analysis and machine learning for the automated evaluation of *Ruscus hypophyllum* foliage. The overall pipeline consists of four main stages: (i) image acquisition under controlled conditions, (ii) image preprocessing and segmentation, (iii) stem size estimation using a geometric

pixel-based method, and (iv) binary classification of foliage condition (good vs. poor) using supervised machine learning models based on Bag of Features representations [23]. Stages (i) and (ii) are shared by both downstream tasks, while stages (iii) and (iv) operate independently on the preprocessed output, allowing each attribute to be assessed using the most appropriate method for its nature. The overall architecture of the proposed system is illustrated in Figure 1.




Figure 1. Architecture of the hybrid framework for *Ruscus hypophyllum* quality assessment, comprising a preprocessing stage followed by parallel geometric and BoF classification branches.

2.1. Image Acquisition

The image dataset was obtained through a controlled acquisition process designed to ensure consistent lighting, scale, and background conditions across all samples. Images of *Ruscus hypophyllum* foliage were captured using a Xiaomi Redmi Note 10 Pro smartphone (Xiaomi Corporation, Beijing, China) equipped with a Samsung ISOCELL HM2 image sensor (Samsung Electronics Co., Ltd., Suwon-si, Gyeonggi-do, Republic of Korea) (maximum resolution: 108 MP), operated through the Open Camera application (v1.53, Android) [25], which allows user-defined output resolutions. The camera was configured to a custom resolution of 3500×1300 pixels, selected to cover the full length of the foliage sample and the reference marker within a single frame while maintaining adequate spatial detail for both geometric measurement and visual feature extraction.

Plant samples were placed horizontally on a flat white surface, with the camera mounted at a fixed distance ranging from 40 to 50 cm above the scene. A reference marker of known physical length was included within the field of view of each image to enable spatial calibration, compensating for any residual variation in the camera-to-sample distance across acquisitions [26]. To improve image quality and minimize shadow formation, the acquisition setup included two white LED lamps of 9 W each, providing approximately 800 lumens per lamp, arranged symmetrically on both sides of the acquisition surface to ensure uniform illumination across all captured samples. Additionally, a white background was used to facilitate the subsequent segmentation of the foliage from the background during image preprocessing.

A total of 1500 images were initially captured. A subsequent manual quality inspection was performed to remove images exhibiting motion blur, partial occlusion of the sample, or incomplete visibility of the reference marker. Following this filtering process, 1233 images were retained to constitute the final dataset used in all experiments.

2.2. Dataset Description

The dataset used in this study comprises 1233 images of *Ruscus hypophyllum* foliage acquired during the post-harvest processing stage under the controlled conditions described in Section 2.1. Each image was annotated with respect to two independent attributes: stem size and foliage condition.

Stem size was categorized into three classes—small, medium, and large—based on the physical length of the main stem, as determined by manual measurement using a reference ruler. The length thresholds defining each category were established according to commercial post-harvest grading standards applied in the production facility [3]: stems up to 25 cm were classified as small, stems between 25 cm and 35 cm as medium, and stems between 35 cm and 45 cm as large.

Foliage condition was defined as a binary attribute—good or poor—based on visual inspection criteria. Samples were labeled as good when they exhibited uniform leaf distribution, adequate leaf density, and absence of visible damage such as yellowing, mechanical injury, or pathological symptoms. Samples were labeled as *poor* when one or more of these criteria were not met. Labeling was performed independently by two post-harvest workers with more than five years of experience in the classification of *Ruscus hypophyllum* foliage at the production facility. Both annotators applied the same visual criteria defined above to assign each sample to the good or poor condition class. Inter-annotator agreement was high, with disagreements occurring in only 5 out of 1233 cases (0.4%). In these cases, the final label was determined by consultation with the facility's senior agronomist, who served as the arbiter. This three-tier annotation protocol ensured consistent and traceable ground truth labels across the full dataset.

Although six combinations of size and condition can theoretically be defined, the two attributes were treated as independent tasks in the proposed framework: stem size was estimated using the geometric method described in Section 2.4, while foliage condition was classified using the machine learning models described in Section 2.6. The distribution of images across both attributes is summarized in Table 2.

Table 2. Dataset distribution by stem size and foliage condition.

Size Class	Length Range (cm)	Images	Condition	Images
Small	$l \leq 25$	407	Good	525
Medium	$25 < l \leq 35$	413	Poor	708
Large	$35 < l \leq 45$	413	—	—
Total	—	1233	Total	1233

The dataset presents a moderate class imbalance in the condition attribute, with 525 good samples (42.6%) and 708 poor samples (57.4%). This imbalance is representative of real post-harvest processing conditions and is addressed during model evaluation through stratified sampling, as described in Section 2.7. The dataset size is consistent with previous studies applying classical machine learning models to structured visual feature representations [19,21], where datasets of similar scale have demonstrated sufficient statistical power for reliable model comparison.

2.3. Image Preprocessing

Prior to feature extraction, all images were subjected to a four-stage preprocessing pipeline designed to improve visual quality, reduce illumination variability, and isolate the foliage region from the background. The entire preprocessing pipeline was implemented in Python 3.11 (Python Software Foundation, Wilmington, DE, USA) using the

scikit-image (version 1.5.2) and OpenCV (version 4.10.0), executed in Google Colaboratory (Google LLC, Mountain View, CA, USA). libraries [27,28]. The pipeline consists of the following stages: resizing, color space conversion, background segmentation, and image enhancement, as illustrated in Figure 2.

Stage 1—Resizing and normalization. Each image was resized so that the full length of the foliage stem was contained within the image frame, ensuring that no relevant morphological information was lost during subsequent processing. No fixed target resolution was imposed, as stem length varies across samples; instead, the resizing factor was determined individually per image to preserve the aspect ratio and guarantee complete stem coverage. Following resizing, intensity normalization was applied to reduce inter-image variability caused by minor illumination differences across acquisition sessions. Intensity normalization is a standard preprocessing step for reducing inter-image illumination variability in agricultural vision systems, where consistent feature representations are critical for reliable classifier performance [15,16].

Stage 2—Color space conversion. Each resized image was converted from the RGB color space to grayscale using a standard luminance-weighted transformation. The grayscale representation was used as the primary input for the subsequent segmentation stage. Additionally, opponent color channels—green-red and blue-yellow—were computed from the original RGB image to support visual inspection and qualitative analysis of foliage chromatic characteristics, including the detection of yellowing and color anomalies associated with poor foliage condition. Grayscale conversion reduces computational complexity while preserving the structural and textural information required for subsequent segmentation and local feature extraction [21].

Stage 3—Background segmentation. Background removal was performed on the grayscale image using Otsu’s global thresholding method [29], which automatically determines an optimal intensity threshold by minimizing the intra-class variance between foreground and background pixels. Given the controlled white background used during image acquisition, this method produced reliable binary masks across all samples. The resulting binary mask was applied to isolate the *Ruscus hypophyllum* foliage from the background in subsequent processing stages, using the `threshold_otsu` function from the `scikit-image` library (version 1.5.2) [27]. Otsu’s thresholding has been widely validated as an effective and parameter-free background removal method in plant image analysis tasks, producing reliable binary masks under controlled white-background acquisition conditions [15,29].

Stage 4—Image enhancement and filtering. To improve the discriminability of local visual features in the segmented foliage region, contrast enhancement was applied using Contrast Limited Adaptive Histogram Equalization (CLAHE) [30], implemented via the `createCLAHE` function from the OpenCV library (version 4.10.0) [28]. CLAHE operates on local image tiles rather than the entire image, enhancing contrast in each region independently while limiting noise amplification in homogeneous areas [30]. Following contrast enhancement, a Gaussian filter was applied using `cv2.GaussianBlur` to reduce high-frequency noise introduced during acquisition and contrast processing, producing a smoothed image suitable for both geometric analysis and local feature extraction. CLAHE has been shown to improve local feature discriminability in plant image classification tasks by enhancing contrast in regions with subtle visual differences [21,30], while Gaussian filtering reduces high-frequency noise that would otherwise interfere with stable keypoint detection in SIFT-based feature extraction pipelines [24]. This enhancement stage was applied uniformly to the output fed into both the geometric estimation branch and the Bag of Features classification branch of the proposed framework.




Figure 2. Image preprocessing pipeline applied to *Ruscus hypophyllum* foliage images. The pipeline consists of four stages: (1) resizing and normalization, (2) color space conversion to grayscale and opponent color channels, (3) background segmentation using Otsu's thresholding, and (4) image enhancement using CLAHE and Gaussian filtering. The output is used as input to both the geometric estimation and the Bag of Features classification branches.

2.4. Stem Size Estimation

Stem size estimation is performed using a geometric multi-stage image processing approach implemented entirely in Python 3.11 using the OpenCV library (version 4.10.0) [28]. By leveraging the controlled acquisition environment, the proposed method enables the conversion of pixel-level measurements into real-world physical dimensions (mm) without relying on machine learning models, following approaches based on geometric calibration and image-based measurement reported in recent agricultural vision systems [16]. The complete pipeline for stem size estimation is illustrated in Figure 3.




Figure 3. Pipeline for stem size estimation based on geometric processing and spatial calibration. The workflow includes bounding box detection (red frame, delimiting the detected plant region), reference marker identification, orientation correction, and pixel-to-millimeter conversion. for stem size estimation based on geometric processing and spatial calibration. The workflow includes bounding box detection, reference marker identification, orientation correction, and pixel-to-millimeter conversion.

2.4.1. Reference Marker Detection

A physical reference marker of known length was included within the field of view of each image during acquisition. This marker consists of a straight line segment of fixed and known physical length, positioned horizontally at 0° with respect to the image frame. The marker was detected automatically in each image using contour detection via the `cv2.findContours` function, followed by blob analysis to identify the marker region based on its geometric properties, including area, aspect ratio, and compactness [28]. Once detected, the marker serves two purposes: (i) it defines the reference orientation of the

coordinate system, since the marker is known to be horizontal ($\theta_{ref} = 0^\circ$), and (ii) it provides the spatial calibration factor S expressed in pixels per millimeter ($\frac{\text{pixels}}{\text{mm}}$), computed as:

$$S = \frac{L_{marker}^{pixels}}{L_{marker}^{mm}} \quad (1)$$

where L_{marker}^{pixels} is the detected length of the marker in pixels and L_{marker}^{mm} is its known physical length in millimeters. This pixel-to-metric conversion approach has been widely applied in precision agriculture and plant phenotyping systems [16,17].

2.4.2. Stem Detection and Orientation Estimation

The main stem structure is identified from the segmented foliage mask by comparison with the reference marker pattern. Since the reference marker is a known horizontal line segment at 0° , the stem region is identified as the primary elongated structure in the binary mask whose orientation deviates from the reference orientation. The inclination angle θ of the stem with respect to the horizontal axis is then estimated by comparing the stem's principal axis against the reference marker direction, using contour-based bounding box analysis via `cv2.minAreaRect` [28]. This angular information is used to perform a deskewing operation, compensating for variations in stem inclination and ensuring a consistent geometric reference frame across all samples.

2.4.3. Trigonometric Length Correction and Size Categorization

After orientation estimation, the stem length is derived from its horizontal projection obtained from the axis-aligned bounding box. Since the stem may be inclined at angle θ with respect to the horizontal axis, a trigonometric correction is applied to recover the true stem length, as commonly used in vision-based measurement systems for inclined objects [31]. The corrected stem length in pixels is computed as:

$$d_{pixels} = \frac{d_x}{\cos(\theta)} \quad (2)$$

where d_x represents the horizontal projection of the stem length obtained from the bounding box, and θ is the inclination angle of the stem with respect to the horizontal axis. The final stem length in millimeters is obtained by applying the spatial calibration factor S derived from Equation (1):

$$L_{mm} = \frac{d_x}{S \cdot \cos(\theta)} \quad (3)$$

The estimated stem length L_{mm} is subsequently used to assign each sample to one of three size categories according to the thresholds defined in Table 2: small ($L_{mm} \leq 250$ mm), medium ($250 < L_{mm} \leq 350$ mm), and large ($350 < L_{mm} \leq 450$ mm).

2.5. Feature Extraction Using Bag of Features

Visual features were extracted using the Bag of Features (BoF) approach [23], which remains a competitive representation for image classification tasks under limited data conditions [22]. Feature extraction plays a critical role in leaf-based classification systems, where characteristics such as shape, texture, and color are commonly used to discriminate between plant samples [21]. Local feature points were detected using the Scale-Invariant Feature Transform (SIFT) detector [24], and descriptors were computed for each keypoint. SIFT was selected due to its well-established robustness to changes in scale, rotation, and illumination conditions, as well as its broad validation in plant-related image classification

tasks [21,24]. The entire feature extraction pipeline was implemented in Python 3.11 using the OpenCV library (version 4.10.0) [28].

The extracted SIFT descriptors were clustered using the k-means algorithm [32], implemented via `sklearn.cluster.KMeans` from the `scikit-learn` library (version 1.5.2) [33], to construct a visual vocabulary of $K = 100$ words. The k-means algorithm was run with a maximum of 300 iterations and a convergence tolerance of 10^{-4} , using ten independent random initializations to reduce sensitivity to the initial centroid placement. To prevent information leakage from test data into the feature representation, the visual vocabulary was constructed independently within each cross-validation fold: in each of the 10 folds, k-means was fitted exclusively on the SIFT descriptors extracted from the training partition (approximately 1110 images), and the resulting vocabulary was subsequently applied to encode the BoF histograms of both the training and test subsets for that fold. Test-fold descriptors were never used during vocabulary construction, ensuring unbiased performance estimates across all folds. The selection of an appropriate vocabulary size is known to directly influence classification performance in BoF-based systems, with excessively small vocabularies producing insufficient visual discrimination and excessively large ones introducing redundancy without accuracy gains [21,23]. The optimal vocabulary size $K = 100$ was determined through a systematic sensitivity analysis described in Section 3.2.

Each image was then represented as a histogram of visual word occurrences:

$$H = (H_1, H_2, \dots, H_K) \quad (4)$$

where $K = 100$ and H_i denotes the frequency of occurrence of the i -th visual word in the image. Prior to classification, the feature vectors were normalized using L2-normalization to ensure unit-length representations, improving numerical stability and classifier performance across models with different sensitivity to feature magnitude [19]. The overall BoF pipeline is illustrated in Figure 4.




Figure 4. Bag of Features extraction pipeline. Local SIFT keypoints are detected and described for each preprocessed image. Descriptors are clustered using k-means to construct a visual vocabulary of $K = 100$ words. Each image is represented as a normalized histogram of visual word occurrences, which serves as the feature vector for the machine learning classifiers. For the 2D projection of the SVM classifier (Step 4), blue points represent the good quality class, while red points represent the poor quality class.

2.6. Machine Learning Classifiers

The classification task was formulated as a binary problem, in which each foliage sample was assigned to one of two condition classes: good or poor. Several supervised machine learning algorithms were evaluated, including Decision Trees, Logistic Regression, Gaussian Naive Bayes, Support Vector Machines (SVM), k-Nearest Neighbors (KNN), and Multilayer Perceptron Neural Networks (MLP) [19,34]. For each family of classifiers, multiple configurations were evaluated by varying key hyperparameters such as kernel type for SVM, number of neighbors for KNN, distance metric, and network architecture for MLP, resulting in a total of 24 evaluated classifier configurations, as detailed in Section 3.

All classifiers were implemented in Python 3.11 using the `scikit-learn` library (version 1.5.2) [33]. To ensure a fair and unbiased comparison between models, default hyperparameter values provided by `scikit-learn` were used for all configurations, with the exception of the structural parameters that define each variant (e.g., kernel type, number of neighbors, hidden layer size). This approach avoids overfitting the model selection to the specific dataset while maintaining reproducibility [19].

The selection of classical machine learning models over deep learning approaches was motivated by three main factors. First, these models are computationally efficient and interpretable, making them suitable for deployment in resource-constrained agricultural processing environments [6]. Second, deep learning approaches such as CNNs typically require substantially larger labeled datasets to achieve reliable generalization—commonly on the order of tens of thousands of samples [8,10]—whereas the present dataset comprises 1233 images, a scale consistent with previous studies applying classical models to structured visual feature representations [19,21]. Third, classical models are particularly well-suited for handling structured feature representations such as normalized BoF histograms, where the dimensionality is moderate and the feature space is explicitly defined [19]. A direct empirical comparison between the proposed BoF-based approach and CNN-based alternatives on this dataset is acknowledged as a limitation of the present work and is identified as a priority direction for future research, as discussed in Section 4.6.

Beyond dataset scale considerations, the choice of SIFT+BoF as a feature representation combined with classical machine learning classifiers offers several advantages that are particularly relevant in the context of industrial post-harvest quality inspection for agricultural exports.

Interpretability and fault traceability. The BoF representation provides an explicit, human-interpretable feature space in which each dimension corresponds to the frequency of a specific visual pattern detected in the image. This transparency allows quality control operators and system engineers to audit classification decisions by inspecting which visual words contributed most to a given prediction, facilitating fault traceability in industrial inspection workflows—a property that end-to-end deep learning models do not inherently provide [6]. Furthermore, classical classifiers produce explicit decision boundaries that can be inspected and validated independently of the feature extraction stage, supporting systematic fault diagnosis in production environments.

Algorithmic stability and predictability. Classical machine learning models exhibit well-characterized and stable behavior under distribution shifts within the trained feature space. Their decision boundaries are deterministic and do not vary across inference runs, which is a critical property for quality certification in export-oriented processing chains where consistent and reproducible grading outcomes are required [3]. This stability applies to the classifier family as a whole and is not contingent on the specific model selected.

Lightweight deployment and maintainability. The proposed pipeline requires no GPU infrastructure and depends exclusively on open-source libraries with stable long-term support (OpenCV, `scikit-learn`, `scikit-image`). Model retraining, when required due

to seasonal or variety-related distribution changes, can be completed rapidly on local hardware without cloud infrastructure or specialized deep learning expertise, reducing operational costs and dependence on external technical support [6]. These properties are shared across all evaluated classical configurations and represent a systematic advantage of the chosen methodological family.

It should be noted that the specific classifier ultimately adopted in this work—the Linear SVM—was not selected a priori as part of the pipeline design, but rather emerged empirically as the best-performing model from the systematic comparative evaluation of 24 configurations across six classical machine learning families reported in Section 3.3, as further supported by the Friedman test and post hoc Wilcoxon signed-rank analysis reported in Section 3.3. The multi-dimensional justification presented above therefore applies at two distinct levels: (i) the a priori choice of SIFT+BoF as the feature representation and of classical machine learning models as the classifier family, and (ii) the a posteriori selection of the Linear SVM specifically, which is grounded in the empirical and statistical results of Section 3.3.

2.7. Performance Evaluation Metrics

The proposed framework is evaluated across two independent tasks, each assessed using a dedicated set of metrics: geometric stem size estimation and binary foliage condition classification.

2.7.1. Geometric Estimation Metrics

The accuracy of the stem size estimation method was evaluated by comparing the estimated stem lengths against ground truth measurements obtained manually using a physical reference ruler. Let L_i^{est} and L_i^{gt} denote the estimated and ground truth stem lengths for the i -th sample, respectively. The following metrics were computed over all N samples:

$$MAE = \frac{1}{N} \sum_{i=1}^N |L_i^{est} - L_i^{gt}| \quad (5)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (L_i^{est} - L_i^{gt})^2} \quad (6)$$

$$\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (e_i - \bar{e})^2} \quad (7)$$

where $e_i = |L_i^{est} - L_i^{gt}|$ and $\bar{e}$ is the mean absolute error. In addition to continuous measurement evaluation, a discrete size categorization accuracy is computed to assess the reliability of the three-class size assignment (small, medium, large):

$$Accuracy_{size} = \frac{N_{correct}}{N_{total}} \quad (8)$$

where $N_{correct}$ is the number of samples correctly assigned to their ground truth size class.

2.7.2. Classification Metrics

The performance of the machine learning classifiers was evaluated using four standard metrics for binary classification: Accuracy, Precision, Recall, and F1-score [35].

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (9)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (10)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (11)$$

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (12)$$

where TP , TN , FP , and FN denote the number of true positives, true negatives, false positives, and false negatives, respectively. Confusion matrices were additionally analyzed to provide a deeper understanding of the classification behavior of each model and to identify systematic misclassification patterns [35].

All classification results were obtained using a stratified 10-fold cross-validation scheme [19]. Stratification ensures that the proportion of good and poor samples is preserved across all folds, which is particularly important given the moderate class imbalance present in the dataset (42.6% good vs. 57.4% poor, as described in Table 2). The use of 10 folds provides a reliable estimate of generalization performance while maintaining a sufficient number of training samples per fold [19].

The complete cross-validation pipeline was executed as follows: in each fold, (i) SIFT descriptors were extracted from the training images; (ii) the visual vocabulary was constructed by fitting k-means exclusively on those training descriptors; (iii) BoF histograms were computed for both training and test images using the fold-specific vocabulary; (iv) histograms were normalized using L2-normalization; and (v) the classifier was trained on the training histograms and evaluated on the test histograms. This procedure guarantees that no information from the test fold influenced either the vocabulary construction or the classifier training at any stage. The reported results correspond to the mean and standard deviation of each metric computed across the 10 folds.

3. Results

This section presents the experimental results obtained for both stages of the proposed framework: geometric stem size estimation and binary foliage condition classification. Results are reported separately for each task, followed by a vocabulary size sensitivity analysis and a comparative confusion matrix analysis of representative classifiers.

3.1. Stem Size Estimation Evaluation

The performance of the proposed geometric method was evaluated by comparing the estimated stem lengths against ground truth measurements obtained manually using a physical reference ruler over all 1233 samples in the dataset. The quantitative results are summarized in Table 3.

Table 3. Quantitative evaluation metrics for the proposed geometric stem size estimation method.

Metric	Value
Mean Absolute Error (MAE)	1.2 mm
Standard Deviation (σ)	0.45 mm
Root Mean Square Error (RMSE)	1.3 mm
Size Classification Accuracy	99.84%
Misclassified Samples	2/1233

The proposed method achieved a MAE of 1.2 mm with a standard deviation of 0.45 mm, and an RMSE of approximately 1.3 mm, confirming the low variability and high consistency of the estimation process across all evaluated samples. The low standard deviation relative

to the MAE indicates that estimation errors are uniformly distributed and do not exhibit systematic bias toward any particular stem size category.

The incorporation of trigonometric orientation correction proved essential for accurate measurement, particularly for inclined stems. Without angular compensation, the horizontal projection of the stem length systematically underestimates the true stem length due to the cosine foreshortening effect described in Equation (2). The use of the reference marker enabled a reliable and consistent conversion from pixel units to real-world dimensions via the spatial calibration factor S defined in Equation (1), ensuring measurement consistency across samples acquired at varying camera-to-sample distances.

Based on the estimated stem length L_{mm} , each sample was assigned to one of three size categories according to the thresholds defined in Section 2.2. Of the 1233 evaluated samples, only 2 were incorrectly categorized, yielding a size classification accuracy of 99.84% (error rate: 0.16%). Both misclassified samples corresponded to stems with lengths near the category boundaries, where measurement uncertainty is inherently higher. These results confirm the robustness and reliability of the proposed geometric approach for accurate and interpretable stem size estimation under controlled acquisition conditions.

3.2. Vocabulary Size Sensitivity Analysis

A systematic sensitivity analysis was conducted to determine the optimal vocabulary size K for the Bag of Features representation. Classification accuracy of the Linear SVM was evaluated across $K \in \{20, 40, \dots, 280, 300\}$ under stratified 10-fold cross-validation. For each value of K , the visual vocabulary was constructed independently within each training fold, following the same leakage-free pipeline described in Section 2.7, ensuring that the reported accuracy values are unbiased estimates of generalization performance for each vocabulary size. As illustrated in Figure 5, accuracy reaches a plateau at $K = 100, 120$, and 140 (92.4%) and remains stable within a narrow range (91.4–91.8%) for larger vocabulary sizes, indicating that further increases in vocabulary size beyond $K = 100$ do not yield appreciable gains in classification performance for this dataset. Values below $K = 100$ show a progressive reduction in accuracy of up to 1.3 percentage points relative to the plateau, suggesting that a minimum vocabulary size is required to adequately capture the visual variability of the *Ruscus hypophyllum* foliage condition classes [21,23]. Based on these results, $K = 100$ was selected as the vocabulary size for all subsequent classification experiments, as it achieves the highest accuracy with the smallest vocabulary size, representing the most favourable trade-off between discriminative capacity and computational cost.




Figure 5. Linear SVM classification accuracy as a function of vocabulary size K under stratified 10-fold cross-validation. Accuracy reaches a plateau at $K = 100$ – 140 (92.4%).

3.3. Classification Performance

Table 4 summarizes the average performance metrics obtained by all evaluated classifiers using stratified 10-fold cross-validation. Classifiers are grouped by family to facilitate comparison across model types.

Table 4. Average performance metrics of evaluated classifiers grouped by family. All results correspond to stratified 10-fold cross-validation. Accuracy and F1-score are reported as mean $\pm$ standard deviation.

No.	Classifier	Acc. (%)	Prec. (%)	Rec. (%)	F1 (%)
<i>Support Vector Machines</i>					
1	Linear SVM	92.4 $\pm$ 1.0	90.7 $\pm$ 1.2	91.4 $\pm$ 1.1	91.1 $\pm$ 1.2
2	Quadratic SVM	91.1 $\pm$ 1.1	89.1 $\pm$ 1.3	90.1 $\pm$ 1.2	89.6 $\pm$ 1.3
3	Cubic SVM	89.1 $\pm$ 1.2	86.7 $\pm$ 1.4	88.0 $\pm$ 1.3	87.3 $\pm$ 1.5
4	Fine Gaussian SVM	81.3 $\pm$ 1.6	77.1 $\pm$ 1.8	80.0 $\pm$ 1.7	78.5 $\pm$ 1.9
5	Mean Gaussian SVM	74.5 $\pm$ 1.9	69.1 $\pm$ 2.1	72.4 $\pm$ 2.0	70.7 $\pm$ 2.3
6	Coarse Gaussian SVM	68.4 $\pm$ 2.1	62.2 $\pm$ 2.3	65.7 $\pm$ 2.2	63.9 $\pm$ 2.5
<i>Neural Networks (MLP)</i>					
7	Medium Neural Network	89.5 $\pm$ 1.5	87.1 $\pm$ 1.6	88.6 $\pm$ 1.7	87.8 $\pm$ 1.8
8	Wide Neural Network	88.6 $\pm$ 1.6	86.3 $\pm$ 1.8	87.2 $\pm$ 1.7	86.7 $\pm$ 1.9
9	Two-layer Neural Network	87.7 $\pm$ 1.6	85.1 $\pm$ 1.8	86.1 $\pm$ 1.7	85.6 $\pm$ 1.9
10	Narrow Neural Network	87.0 $\pm$ 1.8	84.4 $\pm$ 2.0	85.3 $\pm$ 1.9	84.8 $\pm$ 2.1
11	Three-layer Neural Network	84.8 $\pm$ 1.7	81.6 $\pm$ 1.9	82.9 $\pm$ 1.8	82.2 $\pm$ 2.0
<i>Naïve Bayes</i>					
12	Gaussian Naïve Bayes	88.2 $\pm$ 1.7	85.7 $\pm$ 1.9	86.7 $\pm$ 1.8	86.2 $\pm$ 2.0
13	Kernel Naïve Bayes	77.7 $\pm$ 2.0	72.7 $\pm$ 2.2	76.2 $\pm$ 2.1	74.4 $\pm$ 2.4
<i>Logistic Regression</i>					
14	Efficient Logistic Regression	87.3 $\pm$ 1.1	84.7 $\pm$ 1.2	85.7 $\pm$ 1.2	85.2 $\pm$ 1.3
15	Binary Logistic Regression	82.3 $\pm$ 1.3	78.6 $\pm$ 1.5	80.4 $\pm$ 1.4	79.5 $\pm$ 1.6
<i>Decision Trees</i>					
16	Fine Decision Tree	85.8 $\pm$ 2.0	82.4 $\pm$ 2.2	84.8 $\pm$ 2.1	83.6 $\pm$ 2.4
17	Medium Decision Tree	84.2 $\pm$ 2.2	80.6 $\pm$ 2.4	82.9 $\pm$ 2.3	81.7 $\pm$ 2.6
18	Coarse Decision Tree	82.6 $\pm$ 2.4	78.7 $\pm$ 2.6	81.0 $\pm$ 2.5	79.8 $\pm$ 2.8
<i>k-Nearest Neighbors</i>					
19	Cosine KNN	70.0 $\pm$ 2.2	64.5 $\pm$ 2.4	65.7 $\pm$ 2.3	65.1 $\pm$ 2.6
20	Fine KNN	68.0 $\pm$ 2.5	62.0 $\pm$ 2.7	63.8 $\pm$ 2.6	62.9 $\pm$ 2.9
21	Cubic KNN	66.3 $\pm$ 2.6	60.2 $\pm$ 2.8	61.9 $\pm$ 2.7	61.0 $\pm$ 3.0
22	Weighted KNN	64.7 $\pm$ 2.8	58.3 $\pm$ 3.0	60.0 $\pm$ 2.9	59.2 $\pm$ 3.2
23	Medium KNN	63.1 $\pm$ 3.0	56.5 $\pm$ 3.2	58.1 $\pm$ 3.1	57.3 $\pm$ 3.4
24	Coarse KNN	59.9 $\pm$ 3.2	52.7 $\pm$ 3.4	55.2 $\pm$ 3.3	54.0 $\pm$ 3.6

The results reveal a clear performance hierarchy across classifier families. SVM-based models achieved the highest overall performance, with the Linear SVM obtaining the best results across all evaluated metrics: an accuracy of 92.4 $\pm$ 1.0% and an F1-score of 91.1 $\pm$ 1.2%. To provide a statistically rigorous foundation for model selection and account for the overlapping standard deviations among the top-performing configurations, a formal non-parametric statistical analysis was conducted on the fold-by-fold cross-validation accuracy vectors. First, a Friedman test was executed across the 10 folds for the three best-performing configurations overall (Linear SVM, Quadratic SVM, and Medium Neural Network), which confirmed a highly significant performance difference among them (Friedman statistic = 20.00, p = 0.00005). Subsequently, a post hoc pairwise Wilcoxon signed-rank test with Bonferroni correction was performed to compare our top model (Linear SVM) against the runner-up configuration (Quadratic SVM). The post hoc analysis

yielded a corrected p -value of 0.00391 ($p < 0.01$), provides statistical evidence that the superior accuracy of the Linear SVM is statistically significant and due to genuine model efficacy rather than random cross-validation fold variability. The strong performance of linear kernel SVMs is consistent with the characteristics of the BoF feature space, where class boundaries are effectively separable using linear decision surfaces [19,36].

Within the SVM family, a clear degradation in performance is observed as the kernel complexity increases from linear to polynomial (Quadratic, Cubic) and further to Gaussian kernels. This behavior suggests that the BoF histogram feature space does not benefit from higher-order decision boundaries, and that the added flexibility of complex kernels leads to overfitting on this dataset.

MLP Neural Networks constituted the second best-performing family, with accuracies ranging from 84.8% (Three-layer) to 89.5% (Medium), slightly below the top SVM configurations. The performance gap between SVMs and MLPs may be attributed to the limited dataset size, which can restrict the generalization capability of more flexible models requiring larger amounts of training data [10].

Gaussian Naive Bayes achieved a competitive accuracy of 88.2%, outperforming several MLP configurations despite its strong independence assumption, which suggests that the BoF features exhibit a degree of statistical independence consistent with this model's assumptions.

Logistic Regression and Decision Tree models exhibited moderate performance in the range of 82.3–87.3% and 82.6–85.8%, respectively, while KNN configurations consistently yielded the lowest performance across all families, with accuracies ranging from 59.9% to 70.0%. The high variability observed in KNN results (standard deviations up to $\pm 3.6\%$) suggests strong sensitivity to the local feature distribution generated by the BoF representation, where the assumption of proximity in feature space does not effectively capture class-discriminative patterns [19].

3.4. Confusion Matrix Analysis

To provide a deeper understanding of the classification behavior of the evaluated models, a comparative confusion matrix analysis was conducted using three representative classifiers: the best-performing model (Linear SVM), an intermediate model (Medium Neural Network), and the worst-performing model (Coarse KNN). The corresponding confusion matrices are shown in Figure 6.




Figure 6. Confusion matrices for three representative classifiers: (a) Linear SVM, (b) Medium Neural Network, and (c) Coarse KNN. Rows represent ground truth labels and columns represent predicted labels. Values correspond to prediction counts aggregated across all 10 cross-validation folds, totalling 1233 samples.

The confusion matrix of the Linear SVM (Figure 6a) shows a strong concentration of samples along the main diagonal, with 480 true positives (good correctly classified) and

659 true negatives (poor correctly classified) aggregated across all 10 folds, yielding 49 false positives and 45 false negatives in total across the complete dataset of 1233 samples. This reflects a well-balanced trade-off between Precision and Recall (Equations (10) and (11)), consistent with the high F1-score of $91.1 \pm 1.2\%$ reported in Table 4.

The Medium Neural Network (Figure 6b) exhibits a moderate increase in misclassification relative to the Linear SVM, with 69 false positives and 60 false negatives aggregated across all 10 folds, indicating a less optimal decision boundary in the BoF feature space. Despite this, the model maintains an acceptable performance level with an accuracy of $89.5 \pm 1.5\%$, suggesting that MLP-based models can serve as a viable alternative when interpretability requirements are less strict.

In contrast, the Coarse KNN (Figure 6c) shows a substantially higher dispersion of predictions, with 260 false positives and 235 false negatives aggregated across all 10 folds, representing a total misclassification rate of 40.1% across the complete dataset. Notably, this model misclassifies a significant proportion of poor-condition samples as good, resulting in a high false positive rate that would be particularly problematic in a real post-harvest inspection scenario, where undetected poor-quality samples reaching export chains can compromise commercial reputation [3].

From a practical standpoint, the asymmetry between false positives and false negatives has different commercial implications in the post-harvest quality assessment context. A false negative (good sample classified as poor) results in unnecessary product rejection, reducing yield. A false positive (poor sample classified as good) results in defective product reaching the market, with greater commercial and reputational consequences.

4. Discussion

The experimental results obtained for both stages of the proposed framework demonstrate the effectiveness of decoupling geometric measurement from visual classification in the context of post-harvest foliage quality assessment. This section discusses the implications of the obtained results, their consistency with related work, and the limitations of the proposed approach.

4.1. Geometric Stem Size Estimation

The proposed geometric method achieved a MAE of 1.2 mm and an RMSE of approximately 1.3 mm, with a size categorization accuracy of 99.84%. These results confirm that pixel-based geometric measurement combined with spatial calibration and trigonometric orientation correction provides a highly accurate and reliable alternative to learning-based approaches for measurable physical traits under controlled acquisition conditions. This finding is consistent with previous work in plant phenotyping and precision agriculture, where geometric calibration methods have demonstrated competitive accuracy for non-destructive morphological trait estimation [14–16].

The orientation correction step proved particularly relevant for stems with non-negligible inclination angles. Without this correction, the cosine foreshortening effect described in Equation (2) would systematically underestimate stem length, introducing a bias that scales with the inclination angle. The use of a reference marker of known length as both an orientation anchor and a spatial calibration device addresses this challenge elegantly, without requiring additional sensors or complex camera calibration procedures [26]. This makes the proposed geometric stage particularly suitable for deployment in low-cost agricultural processing environments.

The two misclassified samples corresponded to stems with lengths near the category boundaries (i.e., close to 250 mm or 350 mm), where small measurement errors are sufficient to produce an incorrect size assignment. This behavior is inherent to any threshold-based

categorization system and could be mitigated in future work by introducing a rejection region or uncertainty band around the category boundaries, allowing ambiguous samples to be flagged for manual inspection rather than assigned to a potentially incorrect category.

4.2. Foliage Condition Classification

Among the evaluated classifiers, the Linear SVM achieved the best overall performance, with an accuracy of $92.4 \pm 1.0\%$ and an F1-score of $91.1 \pm 1.2\%$. This result is consistent with previous studies in agricultural image classification, where SVM-based models have demonstrated strong effectiveness in high-dimensional structured feature spaces [19,36]. The superiority of the linear kernel over polynomial and Gaussian kernels suggests that the BoF histogram representation produces a feature space where the two classes are approximately linearly separable, making the additional flexibility of non-linear kernels counterproductive due to the risk of overfitting on a dataset of moderate size.

The competitive performance of Gaussian Naive Bayes (88.2%) relative to more complex models such as MLP Neural Networks is noteworthy. This result suggests that the normalized BoF histograms exhibit a degree of feature independence that is compatible with the Naive Bayes assumption, which may be related to the relatively orthogonal nature of visual word occurrences in the vocabulary space [19]. This finding has practical implications, as Gaussian Naive Bayes is computationally inexpensive and highly interpretable, making it a viable lightweight alternative to SVM in resource-constrained deployment scenarios.

The consistently poor performance of KNN configurations (59.9–70.0%) can be attributed to the nature of the BoF feature space. KNN relies on the assumption that samples of the same class are spatially proximate in the feature space. However, BoF histograms are high-dimensional and sparse, producing a feature space where Euclidean and cosine distances may not reliably reflect semantic similarity between samples [19]. This limitation is further evidenced by the high cross-validation variability observed for KNN models (standard deviations up to $\pm 3.6\%$), indicating strong sensitivity to the specific data split used in each fold.

The confusion matrix analysis revealed an important asymmetry in the classification errors of lower-performing models. In the post-harvest context, false negatives (good samples classified as poor) are commercially less consequential than false positives (poor samples classified as good), as they result in unnecessary product rejection rather than defective product reaching export markets [3]. The Linear SVM exhibited the most balanced error distribution between these two types, further supporting its selection as the recommended model for this application.

4.3. Failure Case Analysis

An examination of misclassified samples reveals two systematic error patterns in the Linear SVM output. False negatives (good samples classified as poor, in Figure 7a,b) are predominantly associated with natural leaf surface blemishes and specular reflectance artifacts introduced during image acquisition. These local intensity variations generate BoF visual word distributions that partially overlap with those characteristic of damaged foliage, causing the classifier to incorrectly assign a poor condition label despite the sample meeting all quality criteria. This suggests that the current visual vocabulary may not fully disambiguate between acquisition-induced artifacts and genuine foliar damage at the local feature level.

False positives (poor samples classified as good, in Figure 7c,d) are associated with subtle chlorosis and boundary geometry ambiguity, where early-stage foliar degradation produces chromatic and textural patterns insufficiently distinct from healthy foliage in the BoF representation. These cases represent inherently ambiguous samples that may

also challenge human annotators, as the visual difference between early degradation and acceptable condition is minimal. From a commercial standpoint, false positives carry greater consequence than false negatives in post-harvest inspection, as they result in defective product reaching export markets [3]. Future work could address both error types through targeted data augmentation of boundary-case samples and the incorporation of color-based features explicitly designed to detect early chlorosis patterns.



Figure 7. Representative failure cases of the Linear SVM: false negatives ((a,b), good misclassified as poor) caused by surface blemishes and reflectance artifacts; false positives ((c,d), poor misclassified as good) caused by subtle chlorosis and boundary geometry ambiguity.

4.4. Comparison with Related Work

The proposed hybrid framework achieves competitive results compared to related work in ornamental plant classification. Classification accuracies above 99% for *Anthurium* cultivar identification were reported in [11]; however, that study addressed a multi-class identification problem under highly controlled conditions with a single plant species, whereas the present work targets binary quality grading under the more challenging variability inherent to post-harvest foliage processing. Deep learning-based approaches for ornamental trait estimation have reported accuracies between 35% and 99% depending on the trait [13], highlighting the variability of CNN-based methods across different ornamental attributes and the difficulty of achieving consistent performance without large labeled datasets.

The present work demonstrates that classical machine learning models combined with BoF representations can achieve competitive performance (92.4%) in a practical post-harvest foliage inspection scenario, while offering significant advantages in terms of interpretability, computational efficiency, and reduced data requirements compared to deep learning approaches [8,10]. This positions the proposed framework as a practical and accessible solution for small- and medium-scale ornamental foliage producers seeking to automate quality assessment without the infrastructure requirements of deep learning pipelines.

The competitiveness of classical machine learning models under limited data conditions is supported by empirical evidence in the literature. Comparative studies have consistently shown that SVMs outperform CNNs on small-scale image datasets, while CNNs gain a decisive advantage only when training sets reach the order of tens of thousands of samples [8,10]. Specifically, experimental comparisons on image classification benchmarks have demonstrated that SVM achieves competitive or superior accuracy relative to CNN when the number of training samples is below approximately 1000–2000 instances, a regime in which CNN generalization is constrained by insufficient coverage of the feature space [19,20]. The present dataset of 1233 images falls within this range, providing a principled basis for the selection of classical models. Nevertheless, the question of whether transfer learning with pre-trained CNNs—which can partially compensate for limited labeled data by leveraging features learned from large-scale datasets—would yield competitive or superior results on this specific task remains open and is identified as a priority direction for future work [10].

Notwithstanding the absence of a direct empirical comparison with deep learning baselines, the selection of classical machine learning models in this work is supported by four complementary lines of evidence from the literature.

First, regarding dataset scale and generalization boundary, empirical studies have consistently demonstrated that SVMs achieve competitive or superior accuracy relative to CNNs when the number of training samples is below approximately 1000–2000 instances [19,20]. Too et al. [10] showed that fine-tuned deep learning models for plant identification required datasets of at least 10,000 samples to consistently outperform classical approaches, with performance degrading substantially under limited data conditions even with transfer learning. The present dataset of 1233 images falls within the regime where classical models are empirically competitive.

Second, regarding domain shift in transfer learning, ImageNet-pretrained features are optimized for object recognition in natural scene photographs, whereas the present task involves fine-grained texture discrimination in controlled top-view foliage images with uniform backgrounds. This domain gap limits the direct transferability of pre-trained features and increases the risk of overfitting during fine-tuning on a dataset of moderate size [8,10].

Third, regarding feature space suitability, the BoF representation produces a normalized histogram of fixed dimensionality ($K = 100$) that explicitly encodes local texture patterns relevant for distinguishing good from poor foliage condition. Linear SVMs are theoretically well-suited to this type of structured, moderate-dimensionality feature space [19], a property confirmed by the empirical results obtained in this work.

Fourth, regarding industrial deployment constraints, the target application context—small- and medium-scale post-harvest processing facilities—favors lightweight, interpretable, and maintainable systems over architecturally complex deep learning pipelines that require GPU infrastructure and specialized technical personnel [6]. These practical constraints constitute an independent motivation for the methodological choices made in this work, complementary to the dataset scale argument.

4.5. Practical Deployment Considerations

The proposed framework was evaluated on a standard CPU-based environment (x 86_64 architecture, 2 logical cores, 12.67 GB RAM) to assess its suitability for deployment in resource-constrained agricultural processing facilities. The complete inference pipeline—encompassing image preprocessing, BoF feature extraction, and SVM classification—achieves a mean processing time of 84.2 ms per sample, corresponding to a throughput of approximately 11.9 images per second. These results confirm that the proposed system operates in near-real-time on commodity hardware without requiring dedicated GPU infrastructure, making it practically accessible for small- and medium-scale ornamental foliage producers. The use of open-source Python 3.11 libraries (scikit-learn (version 1.5.2), OpenCV, scikit-image (version 4.10.0)) further reduces deployment costs, as no commercial software licenses are required. The system can be deployed on a standard laptop or low-cost single-board computer connected to a consumer-grade smartphone camera, consistent with the low-cost acquisition setup described in Section 2.1.

Regarding framework transferability, the proposed pipeline is applicable to other single-stem, multi-leaf ornamental foliage species with similar morphological structure, such as *Ruscus aculeatus* or *Asparagus setaceus*. The geometric measurement stage requires only recalibration of the size category thresholds to match the commercial grading standards of the target species, with no algorithmic modifications. The BoF-SVM classification stage requires retraining on a species-specific labeled dataset, as both the visual vocabulary—defined by the k-means centroids fitted on SIFT descriptors—and the SVM decision boundary are learned from the training data and are therefore species-dependent. The migration cost is moderate: a new annotated dataset of comparable scale (approximately 1000–1500 images) and a retraining procedure that can be completed in under one

hour on standard CPU hardware are sufficient, with no architectural modifications to the pipeline. The optimal migration strategy consists of three steps: (i) acquire and annotate a species-specific dataset following the same controlled acquisition protocol described in Section 2.1; (ii) retrain the BoF vocabulary and SVM classifier using the fold-internal pipeline described in Section 2.7; and (iii) validate performance using the same stratified 10-fold cross-validation scheme, reporting the four metrics defined in Section 2.7 to enable direct comparison with the results obtained for *Ruscus hypophyllum*.

4.6. Limitations

Despite the promising results, several limitations of the proposed framework should be acknowledged. First, all experiments were conducted under controlled acquisition conditions, including uniform white background, fixed illumination, and a dedicated reference marker. Performance may degrade under more variable real-world conditions, such as non-uniform backgrounds, varying lighting, or occlusions, which are common in field processing environments.

Addressing this limitation will require targeted data augmentation strategies designed to simulate uncontrolled acquisition conditions, including random brightness and contrast adjustment ($\pm 30\%$), Gaussian noise injection ($\sigma \in [0.01, 0.05]$), image affine transformations (rotation $\pm 15^\circ$, shear $\pm 10\%$), and synthetic background replacement using randomly sampled textures. On the hardware side, integration into existing post-harvest conveyor systems would benefit from an industrial line-scan camera with automatic exposure control, active LED ring illumination with diffuser panels to minimize specular reflections, and a standardized enclosure to isolate the acquisition zone from ambient lighting variability. System generalization will be evaluated following a phased incremental protocol progressing from (i) controlled laboratory conditions (current baseline), through (ii) semi-controlled conditions with variable illumination and non-uniform backgrounds, to (iii) real-world field processing environments with dynamic backgrounds and uncontrolled lighting, enabling direct performance comparison across conditions using the same stratified 10-fold cross-validation protocol applied in the present work.

Second, the dataset size of 1233 images, while sufficient for the evaluated classical models, remains relatively limited for more data-intensive approaches and may restrict the generalization of the results to broader acquisition conditions. Third, the foliage condition labels were assigned through a three-tier annotation protocol; future work should compute formal inter-annotator agreement metrics (e.g., Cohen's κ) to further quantify label consistency. Fourth, the current framework does not incorporate a rejection mechanism for ambiguous samples near classification or size boundaries, which could improve the reliability of the system in practical deployment scenarios. Finally, the present work does not include a direct empirical comparison with CNN-based approaches on the same dataset. Although the dataset size of 1233 images is consistent with the scale typically associated with classical machine learning models [10,19], the question of whether transfer learning with pre-trained CNNs—which can partially compensate for limited data—would yield competitive or superior results on this specific task remains open. This comparison is identified as a priority direction for future work.

Regarding reference marker detection robustness, the current implementation requires successful detection of the reference marker in every image to perform spatial calibration. In the event of marker occlusion, soiling, or excessive specular reflection leading to detection failure, the system flags the affected image for manual inspection rather than proceeding with an unreliable calibration estimate. As a contingency for partial marker visibility, a fallback calibration scheme based on a fixed nominal scale factor—derived from the known mean camera-to-sample distance and the observed distribution of calibration factors S

across the dataset—could be applied with an associated uncertainty estimate. This fallback mechanism has not been implemented in the current version but is identified as a practical robustness improvement for future deployment, particularly relevant for integration into conveyor-based processing lines where marker visibility cannot always be guaranteed.

5. Conclusions

This paper presented a two-stage hybrid computer vision framework for the automated quality assessment of *Ruscus hypophyllum* foliage in post-harvest processing environments. The proposed system explicitly decouples geometric stem size estimation from visual foliage condition classification, addressing a gap identified in the literature where heterogeneous quality attributes are typically treated within a single unified learning framework. This decoupling improves system interpretability, reduces model complexity, and allows each attribute to be assessed using the most appropriate method for its nature.

The geometric stem size estimation method, based on spatial calibration via a reference marker and trigonometric orientation correction, achieved a Mean Absolute Error of 1.2 mm, an RMSE of approximately 1.3 mm, and a size categorization accuracy of 99.84% over 1233 evaluated samples. These results demonstrate that deterministic geometric methods can provide highly accurate and reproducible measurements of physical plant traits under controlled acquisition conditions, without requiring any learning-based component for this task.

For foliage condition classification, the Bag of Features representation extracted using SIFT descriptors, combined with a Linear SVM classifier, achieved an accuracy of $92.4 \pm 1.0\%$ and an F1-score of $91.1 \pm 1.2\%$ using stratified 10-fold cross-validation. The comparative evaluation of 24 classifier configurations across six families demonstrated that the linear separability of the BoF feature space favors SVM-based models, while KNN configurations were systematically outperformed due to their sensitivity to the high-dimensional sparse nature of histogram representations. These findings provide evidence-based guidance for model selection in structured visual feature spaces for agricultural inspection tasks.

From a practical perspective, the proposed framework offers a low-cost, interpretable, and computationally efficient solution for automated foliage quality grading, with an inference time of 84.2 ms per sample on standard CPU hardware, suitable for deployment in small- and medium-scale ornamental foliage processing facilities. The use of a consumer-grade smartphone camera and open-source Python 3.11 libraries reduces the infrastructure requirements for adoption, making the system accessible to producers in export-oriented agricultural sectors such as the Ecuadorian ornamental foliage industry [2].

Several promising avenues for future research emerge from this study to transition the proposed methodology toward field-ready deployment.

First, to establish robust generalization under real-world processing environments, future efforts will evaluate system performance following a structured three-phase incremental protocol: (i) laboratory-scale replication under controlled changes, (ii) semi-controlled conditions featuring variable illumination and non-uniform backgrounds, and (iii) fully uncontrolled real-world field environments. To support this progression, essential hardware adaptations for on-site deployment will be integrated, including an industrial line-scan camera equipped with automatic exposure control, an active LED ring illumination system with diffuser panels, and a standardized enclosure designed to isolate the acquisition zone from ambient lighting variability.

Second, a direct empirical comparison between classical computer vision techniques and end-to-end deep learning approaches will be conducted. Specifically, convolutional neural networks (CNNs) such as EfficientNet-B0 and MobileNetV2 will be trained and

evaluated on the same dataset. This architectural comparison will follow the exact stratified 10-fold cross-validation protocol applied in the present work to ensure an objective performance baseline across both frameworks [8].

Third, to address the data scarcity inherent to specialized agricultural domains and mitigate potential overfitting in deep learning models, a rigorous data augmentation strategy will be deployed [10]. This framework will encompass localized geometric transformations (e.g., random rotations within $\pm 15^\circ$) and radiometric perturbations, including Gaussian noise injection ($\sigma \in [0.01, 0.05]$) and brightness adjustments between -20% and $+20\%$, artificially expanding the variance of the training set.

Fourth, future versions of the system will incorporate a confidence-based rejection mechanism within the classification pipeline. By defining strict decision thresholds, the classifier will automatically flag low-confidence predictions, routing ambiguous or heavily degraded fragments to manual expert inspection, thereby maintaining high operational accuracy and preventing false positives in critical agricultural diagnostic workflows.

Author Contributions: F.O.-L.: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing—original draft, Visualization, and overall coordination of all stages of the manuscript (lead author). F.T.-R.: Writing—review and editing, with emphasis on manuscript drafting and language refinement. D.P.: Image preparation, Data curation, Investigation, and Validation. M.H.S.: Methodological review and Writing—review and editing. D.H.P.-O.: Methodological review, Supervision, Project administration, and Writing—review and editing (style and structure). All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Universidad de las Fuerzas Armadas ESPE.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The complete image dataset supporting the conclusions of this article is publicly available in Zenodo at <https://zenodo.org/records/20800717> (accessed on 21 June 2026).

Acknowledgments: The authors would like to thank the Laboratory of Robotics and Industrial Automation at the Universidad de las Fuerzas Armadas ESPE for the support provided during this research. As well, authors acknowledge the valuable support given by the SDAS Research Group (<https://sdas-group.com/>). During the preparation of this manuscript, the authors used large language model (LLM, Claude Sonnet 5, Anthropic)-based writing assistance tools to support language editing, text structuring, and drafting of selected sections. The authors have reviewed, edited, and take full responsibility for all content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BoF	Bag of Features
CLAHE	Contrast Limited Adaptive Histogram Equalization
CNN	Convolutional Neural Network
KNN	k-Nearest Neighbors
MAE	Mean Absolute Error
MLP	Multilayer Perceptron
RMSE	Root Mean Square Error
SIFT	Scale-Invariant Feature Transform
SVM	Support Vector Machine

References

- Meir, S.; Philosoph-Hadas, S. Postharvest Physiology of Ornamentals: Processes and Their Regulation. *Agronomy* **2021**, *11*, 2387. [CrossRef]
- Guaita-Pradas, I.; Rodríguez-Mañay, L.O.; Marques-Perez, I. Competitiveness of Ecuador's Flower Industry in the Global Market in the Period 2016–2020. *Sustainability* **2023**, *15*, 5821. [CrossRef]
- Franzoni, G.; Rovera, C.; Farris, S.; Ferrante, A. Packaging materials and their effect on ruscus quality changes during storage and vase life. *Postharvest Biol. Technol.* **2024**, *210*, 112789. [CrossRef]
- Ahmad, I.; Saeed, H.A.u.R.; Khan, M.A.S. Ornamental Horticulture: Economic Importance, Current Scenario and Future Prospects. In *Etiology and Integrated Management of Economically Important Fungal Diseases of Ornamental Palms*; Sustainability in Plant and Crop Protection; Ul Haq, I., Ijaz, S., Eds.; Springer: Cham, Switzerland, 2020; Volume 16, pp. 3–40. [CrossRef]
- Cubero, S.; Aleixos, N.; Moltó, E.; Gómez-Sanchis, J.; Blasco, J. Advances in machine vision applications for automatic inspection and quality evaluation of fruits and vegetables. *Food Bioprocess Technol.* **2011**, *4*, 487–504. [CrossRef]
- Liakos, K.G.; Busato, P.; Moshou, D.; Pearson, S.; Bochtis, D. Machine Learning in Agriculture: A Review. *Sensors* **2018**, *18*, 2674. [CrossRef] [PubMed]
- Blasco, J.; Munera, S.; Aleixos, N.; Cubero, S.; Moltó, E. Machine Vision-Based Measurement Systems for Fruit and Vegetable Quality Control in Postharvest. *Adv. Biochem. Eng./Biotechnol.* **2017**, *161*, 71–91. [CrossRef] [PubMed]
- Kamilaris, A.; Prenafeta-Boldú, F.X. Deep learning in agriculture: A survey. *Comput. Electron. Agric.* **2018**, *147*, 70–90. [CrossRef]
- Mohanty, S.P.; Hughes, D.P.; Salathé, M. Using deep learning for image-based plant disease detection. *Front. Plant Sci.* **2016**, *7*, 1419. [CrossRef] [PubMed]
- Too, E.C.; Yujian, L.; Njuki, S.; Yingchun, L. A comparative study of fine-tuning deep learning models for plant disease identification. *Comput. Electron. Agric.* **2019**, *161*, 272–279. [CrossRef]
- Soleimanipour, A.; Chegini, G.R. A Vision-Based Hybrid Approach for Identification of *Anthurium* Flower Cultivars. *Comput. Electron. Agric.* **2020**, *174*, 105460. [CrossRef]
- Sun, X.; Li, Z.; Zhu, T.; Ni, C. Four-Dimension Deep Learning Method for Flower Quality Grading with Depth Information. *Electronics* **2021**, *10*, 2353. [CrossRef]
- Afonso, M.; Paulo, M.J.; Fonteijn, H.; van den Helder, M.; Zwinkels, H.; Rijsbergen, M.; van Hameren, G.; Haegens, R.W. Automatic trait estimation in floriculture using computer vision and deep learning. *Smart Agric. Technol.* **2024**, *7*, 100383. [CrossRef]
- Minervini, M.; Scharr, H.; Tsaftaris, S.A. Image Analysis: The New Bottleneck in Plant Phenotyping. *IEEE Signal Process. Mag.* **2015**, *32*, 126–131. [CrossRef]
- Meraj, T.; Sharif, M.I.; Raza, M.; Alabrah, A.; Kadry, S.; Gandomi, A.H. Computer Vision-Based Plants Phenotyping: A Comprehensive Survey. *iScience* **2023**, *26*, 108709. [CrossRef] [PubMed]
- Putra, B.T.W.; Ramadhani, N.J.; Soedibyo, D.W.; Marhaenanto, B.; Indarto, I.; Yualianto, Y. The use of computer vision to estimate tree diameter and circumference in homogeneous and production forests using a non-contact method. *For. Sci. Technol.* **2021**, *17*, 32–38. [CrossRef]
- Nguyen, T.T.; Vandevorode, J.; Saeys, W. Measurement of Plant Traits Using Computer Vision: A Review of Calibration and Scaling Methods. *Bot. Rev.* **2024**, *90*, 251–311. [CrossRef]
- Franco, V.R.; Hott, M.C.; Andrade, R.G.; Goliatt, L. Hybrid machine learning methods combined with computer vision approaches to estimate biophysical parameters of pastures. *Evol. Intell.* **2023**, *16*, 1271–1284. [CrossRef]
- Kotsiantis, S.B. Supervised machine learning: A review of classification techniques. *Informatica* **2007**, *31*, 249–268.
- Hastie, T.; Tibshirani, R.; Friedman, J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed.; Springer: New York, NY, USA, 2009. [CrossRef]
- Azlah, M.A.F.; Chua, L.S.; Rahmad, F.R.; Abdullah, F.I.; Wan Alwi, S.R. Review on Techniques for Plant Leaf Classification and Recognition. *Computers* **2019**, *8*, 77. [CrossRef]
- Sultana, F.; Sufian, A.; Dutta, P. Advancements in Image Classification Using Convolutional Neural Networks: A Survey. In *Intelligent Computing: Image Processing Based Applications*; Advances in Intelligent Systems and Computing; Springer: Singapore, 2020; Volume 1157, pp. 1–16. [CrossRef]
- Csurka, G.; Dance, C.; Fan, L.; Willamowski, J.; Bray, C. Visual categorization with bags of keypoints. In Proceedings of the ECCV Workshop on Statistical Learning in Computer Vision, Prague, Czech Republic, 15 May 2004; pp. 1–22.
- Lowe, D.G. Distinctive Image Features from Scale-Invariant Keypoints. *Int. J. Comput. Vis.* **2004**, *60*, 91–110. [CrossRef]
- Brewster, M. Open Camera. Android Application, Version 1.53. 2021. Available online: <https://opencamera.org.uk> (accessed on 15 March 2024).
- Chen, H.; Zhuang, J.; Liu, B.; Wang, L.; Zhang, L. Camera calibration method based on circular array calibration board. *Syst. Sci. Control Eng.* **2023**, *11*, 2163914. [CrossRef]

27. van der Walt, S.; Schönberger, J.L.; Nunez-Iglesias, J.; Boulogne, F.; Warner, J.D.; Yager, N.; Gouillart, E.; Yu, Y. scikit-image: Image Processing in Python. *PeerJ* **2014**, *2*, e453. [[CrossRef](#)] [[PubMed](#)]
28. Bradski, G.; Kaehler, A. *Learning OpenCV: Computer Vision with the OpenCV Library*; O'Reilly Media: Sebastopol, CA, USA, 2008.
29. Otsu, N. A Threshold Selection Method from Gray-Level Histograms. *IEEE Trans. Syst. Man Cybern.* **1979**, *9*, 62–66. [[CrossRef](#)]
30. Zuiderveld, K. Contrast Limited Adaptive Histogram Equalization. In *Graphics Gems IV*; Heckbert, P.S., Ed.; Academic Press: San Diego, CA, USA, 1994; pp. 474–485. [[CrossRef](#)]
31. Liu, T.; Lv, X.; Jin, M. Research on Image Measurement Method of Flat Parts Based on the Adaptive Chord Inclination Angle Algorithm. *Appl. Sci.* **2023**, *13*, 1641. [[CrossRef](#)]
32. Lloyd, S.P. Least Squares Quantization in PCM. *IEEE Trans. Inf. Theory* **1982**, *28*, 129–137. [[CrossRef](#)]
33. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
34. Alshammari, A.H.; Bencsik, G.; Ali, H.A. A Survey of Six Classical Classifiers. *Algorithms* **2026**, *19*, 37. [[CrossRef](#)]
35. Powers, D.M.W. Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness and Correlation. *J. Mach. Learn. Technol.* **2020**, *11*, 37–63. [[CrossRef](#)]
36. Ahmad, I.; Alqahtani, A.; Al-Makhadmeh, Z.; Al-Amri, A. Improved Support Vector Machine Enabled Radial Basis Function and Linear Variants for Remote Sensing Image Classification. *Sensors* **2021**, *21*, 4431. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.