

Feature extraction for plant growth estimation

SA Ngorima^{1,2,3}[0000-0002-0775-3529], ASJ Helberg¹[0000-0001-6833-5163], and
MH Davel^{1,2,3}[0000-0003-3103-5858]

¹ Faculty of Engineering, North-West University, South Africa

² Centre for Artificial Intelligence Research, South Africa

³ National Institute for Theoretical and Computational Sciences, South Africa

`aldringorima@gmail.com`

`albert.helberg@nwu.ac.za`

`marelie.davel@nwu.ac.za`

Abstract. Precision agriculture requires the estimation of plant growth stages in real-time. When the plant growth stage is known, the wastage of resources in cultivation, such as nutrients and water, is reduced as only the required resources need to be supplied. Plants at different growth stages, however, have similar morphological features, which can make autonomous growth stage estimation difficult. This paper presents two feature extraction methods for growth stage estimation: one that uses a bank of Gabor filters and morphological operations, and the other that uses pre-trained convolutional neural networks (CNNs) and transfer learning. We test these methods on a publicly available plant growth stage dataset (“bccr-segset”) for two species, canola and radish, grown and captured under indoor conditions. The two proposed feature extraction methods are compared, using support vector machines and boosted trees as classifiers. We find that both methods are suitable for real-time applications, and that CNN features outperform the hand-crafted features, both with regard to speed and accuracy. The best system (VGG-19 features, classified with a radial basis function support vector machine) obtained an accuracy of 98.4% for both species, processing an image in 0.08 seconds.

Keywords: growth stage estimation · precision agriculture · transfer learning · convolutional neural networks · Gabor filters · morphological operations · support vector machines.

1 Introduction

Indoor farming systems, such as vertical farming, are well known for their ability to produce large yields while using fewer resources and taking up less space [26]. In addition, they promote clean food production, as the use of chemicals such as herbicides is not typically required, since farming is done indoors and pests can often be controlled manually. Reducing the use of herbicides also reduces the number of farming resources required. Indoor farming also enables extensive data collection on plants, which is challenging to achieve in conventional farming.

Precision agriculture (PA), which involves the use of technology and analytical methods in agriculture, has significantly improved farming management and crop yields [13, 25]. Precision agriculture in indoor environments utilises cameras, sensors, and actuators to identify and optimise the optimal ranges of the most critical environmental parameters for plant health and growth [25].

Strengthening PA in indoor farming requires accurate plant growth stage estimates to identify the plant development stage. The plant development stage helps determine the appropriate nutrients for plants and it can also be used to predict yield [10]. Applying only necessary resources to crops reduces farming costs. Broad-leaved plants in indoor agriculture grow quickly but are sensitive and require precise nutrition and controlled environments for high yield and quality. Good results are achieved when growth stage estimation is done regularly and preferably in real-time [2]. Several studies have been conducted on plant classification, achieving good results [10, 12, 28, 11, 13]. To the best of our knowledge, the estimation of the growth stage of indoor-grown broad-leaved plants has not yet been thoroughly investigated.

We propose a computer vision-based approach for indoor farming that integrates feature extraction, morphological operations, and machine learning classification to estimate plant growth stages. We develop and compare two feature extraction methods as input to a real-time classifier. The first approach extracts features using Gabor filters, which mainly focus on textural features. The second approach uses CNN-based transfer learning for feature extraction, combined with classical machine learning classifiers. While transfer learning is not new to plant classification, its application for growth stage estimation remains limited. Side-by-side experiments are conducted and the results are compared using accuracy, confusion matrices, and execution time as metrics.

The main contributions of this paper are:

- We show that a pre-trained CNN model that uses transfer learning for plant feature extraction, concatenated with a classical machine learning algorithm, can yield accurate growth stage estimation, even when using a small training dataset. This method is fast and capable of predicting events in real-time.
- Gabor filters are presented as a plant feature extractor that can also produce good classification results. We demonstrate that morphological operations can enhance the Gabor filter’s feature extraction method, as they precisely highlight the shape and edges of an object within the frame. Manually designing a feature extraction method does not require large datasets.

The paper is organised as follows: Section 2 reviews related work on plant detection and growth stage classification. Section 3 describes the background and development of both the Gabor method and the transfer learning-based approach. In Section 4, we describe the dataset, including its preprocessing and preparation to fit the task. Section 5 outlines the classification algorithms (RBF-SVM and XGBoost) and their implementation. Finally, Section 6 presents the results, along with some additional analysis.

2 Related Work

Machine learning techniques have transformed precision agriculture by increasing productivity while reducing costs [29, 13, 25]. Computer vision systems enable visual monitoring of plants, providing comprehensive insights into plant growth and health [16]. However, real-world applications face significant challenges: complex backgrounds, inconsistent lighting, and occlusion can make image segmentation and object recognition difficult [15].

Plant classification typically relies on leaf features such as shape, colour, and texture. Waldchen et al. [28] reviewed 120 studies and found that leaf shape and colour were the most widely used, with only 24 studies focusing on texture. This seems like a missed opportunity. The texture of the plant leaf changes significantly during the growth stages due to anatomical developments: changes in the epidermis, tissue reorganisation, and changes in the venation pattern. These changes may be most visible through texture variations, suggesting that texture-based approaches could be more discriminative for growth-stage classification. Gabor filters have shown superior performance in texture analysis in various applications [5, 22, 27]. Their application to plant growth stage classification has not been explored.

Recent growth stage classification studies have used various approaches. CNN-based methods have shown promise, with Murata et al. [18] achieving an accuracy of (89.3%) for the estimation of the rice growth stage, and Rasti et al. [19] successfully identified 12 growth stages of wheat and 11 growth stages of barley. However, these approaches typically require large datasets and have a prolonged inference time [11]. Colour-based methods have also shown effectiveness, with the latest studies achieving 79-88% accuracy for lettuce growth stage classification [4].

Transfer learning has emerged as an effective approach for plant classification tasks, adapting pre-trained CNN knowledge to new domains while reducing computational requirements and training data needs [13]. Popular architectures include VGG [24], and ResNet [9], typically pre-trained on ImageNet. However, whether features learnt from general object recognition translate effectively to the task of distinguishing consecutive growth stages, where morphological differences are often subtle, remains unclear.

3 Method

This study proposes two feature extraction approaches for plant growth stage estimation: a Gabor filter method enhanced with morphological operations and a CNN transfer learning approach. The Gabor method targets textural features while leveraging morphological operators to highlight leaf contour, shape, and size characteristics. The transfer learning method uses pre-trained convolutional neural networks to extract plant features: specifically, the VGG-16, VGG-19, and ResNet-50 models, chosen for their demonstrated performance in various domains.

3.1 Gabor Filter Method

A Gabor filter [8] combines a sinusoidal wave with a Gaussian envelope to analyse the texture of the image by locating the frequency content in specified directions. The Gaussian envelope captures spatial details, while the sinusoidal component identifies frequency-specific information [7]. This combination appears particularly suited for plant analysis as it remains relatively stable under varying illumination and contrast conditions.

The filtering process involves convolving the image with multiple Gabor kernels at varying frequencies and orientations. The parameters of each kernel determine the type and scale of the features extracted. Specifically, the sinusoidal wave’s frequency (λ) and orientation (θ) define the type of texture captured, while the Gaussian envelope’s standard deviation (σ) and aspect ratio (γ) control the spatial extent and shape of the feature analysis region. Mathematically, the general complex Gabor filter is defined as:

$$g(x, y, \theta, \lambda, \psi, \sigma, \gamma) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \exp\left(i\left(2\pi\frac{x'}{\lambda} + \psi\right)\right) \quad (1)$$

with real and imaginary components given respectively by:

$$g_R(x, y, \theta, \lambda, \psi, \sigma, \gamma) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \cos\left(2\pi\frac{x'}{\lambda} + \psi\right) \quad (2)$$

$$g_I(x, y, \theta, \lambda, \psi, \sigma, \gamma) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \sin\left(2\pi\frac{x'}{\lambda} + \psi\right) \quad (3)$$

where $x' = x \cos(\theta) + y \sin(\theta)$ and $y' = y \cos(\theta) - x \sin(\theta)$. The parameters control the filter behaviour: θ represents orientation, λ defines wavelength, σ sets the Gaussian standard deviation, ψ adjusts phase offset, and γ determines the spatial aspect ratio.

Parameter Selection and Implementation: We create a Gabor filter bank by empirically determining optimal parameter combinations through testing on representative training samples from each class. Parameter values are selected by examining PCA-t-SNE visualisations of extracted features and the monitoring validation score. This process is time-consuming but essential for capturing meaningful textural distinctions specific to plant images.

The phase offset had a minimal impact on classification accuracy, so we set $\psi = 0$ for simplicity. Table 1 shows the final parameter selection after extensive experimentation. The filter bank uses four distinct frequencies across 8 orientations, generating 32 distinct response maps per image.

The 3×3 kernel size was retained while we extensively tuned other Gabor parameters. Rimiru et al. [21] found that a 3×3 Gabor kernel achieved 99.68% efficiency on CIFAR-10, compared to 98.69% and 98.46% for 5×5 and 7×7 kernels, respectively, suggesting diminishing returns from larger kernels in their architecture. We adopt the 3×3 kernel based on this evidence of diminishing returns from larger kernels and its computational efficiency for real-time performance.

Table 1. Gabor filter parameters.

| Filter | Kernel size | θ | λ | ψ | σ | γ |
|--------|--------------|-----------------|------------------|--------|----------|----------|
| 1 | 3×3 | $\frac{\pi}{4}$ | $\frac{\pi}{10}$ | 0 | 0.5 | 0.5 |
| 2 | 3×3 | $\frac{\pi}{4}$ | $\frac{\pi}{8}$ | 0 | 0.5 | 0.5 |
| 3 | 3×3 | $\frac{\pi}{2}$ | $\frac{\pi}{10}$ | 0 | 0.5 | 0.5 |
| 4 | 3×3 | $\frac{\pi}{2}$ | $\frac{\pi}{8}$ | 0 | 0.5 | 0.5 |

Morphological Operations and Feature Extraction Workflow: Figure 1 shows the complete feature extraction workflow using Gabor filters combined with morphological preprocessing. Initially, images are resized to 128×128 pixels and converted to greyscale. Morphological operations are then applied to highlight the contours and shape characteristics of plants. This preprocessing involves opening (noise removal), closing (object enhancement), thresholding (plant segmentation), and contour masking (shape isolation), as illustrated step-by-step in Figures 2a through 2e.

Both original greyscale and morphologically processed contour-masked images undergo convolution with the Gabor filter bank. This generates multiple feature response maps, as demonstrated in Figure 3, where each filter response highlights specific textures and structural plant features. These feature maps subsequently form the input for the machine learning classifiers.

3.2 Transfer Learning Method

Pre-trained CNNs offer an alternative feature extraction approach by leveraging knowledge gained from large-scale datasets. CNN architectures typically stack convolutional layers for feature extraction, pooling layers for dimensionality reduction, and fully-connected layers for classification [11]. We repurpose these networks by removing classification layers and using only the convolutional components as feature extractors. This approach seems promising because convolutional layers capture general visual elements, edges, colours, and textures, which may transfer well to plant analysis [17]. Our implementation uses the pre-trained convolutional layers of VGG-16, VGG-19 [24], and ResNet-50 [9] without retraining these components.

Images are first greyscaled and fed directly to these feature extractors. We deliberately avoid fully-connected layers since they contain most network parameters; around 123 million of VGG-16’s 138 million parameters are found in classification layers [20]. Replacing these computationally expensive components with classical machine learning algorithms is expected to improve prediction speed while maintaining feature quality, which aligns with our real-time processing goals.

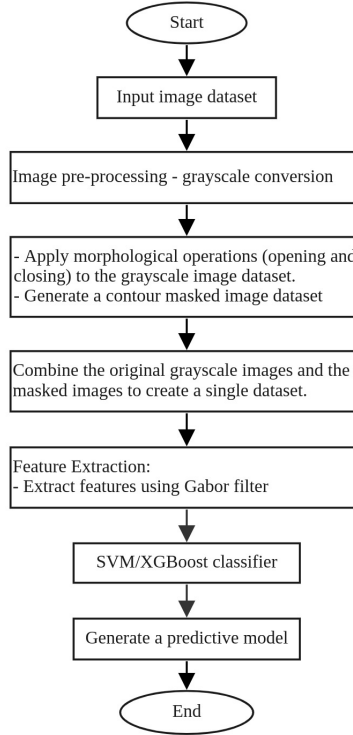


Fig. 1. Gabor method flowchart with morphological operators.

4 Dataset

The dataset used in this work is the publicly available “bccr-segset”, originally compiled by Edith Cowan University, Australia, for crop and weed detection [12]. It contains approximately 30,000 images of canola, radish, and maize captured at four distinct growth stages under controlled indoor conditions, with images acquired at 228×228 pixels from a top-mounted camera. For our study, we repurposed this data set for the classification of growth stages, focusing exclusively on broad-leafed plants (canola and radish). Maize images were excluded, and a separate “background” class was added to remove bias from empty frames. To address class imbalance, we used stratified sampling during dataset partitioning and stratified cross-validation for evaluation.

4.1 Dataset partitioning

We grouped images according to the pre-labelled growth stage and plant species. Each dataset was partitioned using stratified random sampling to maintain proportional class representation across all growth stages: approximately 90% for

Feature extraction for plant growth estimation




Fig. 2. Morphological operations workflow: original to contour-masked image.

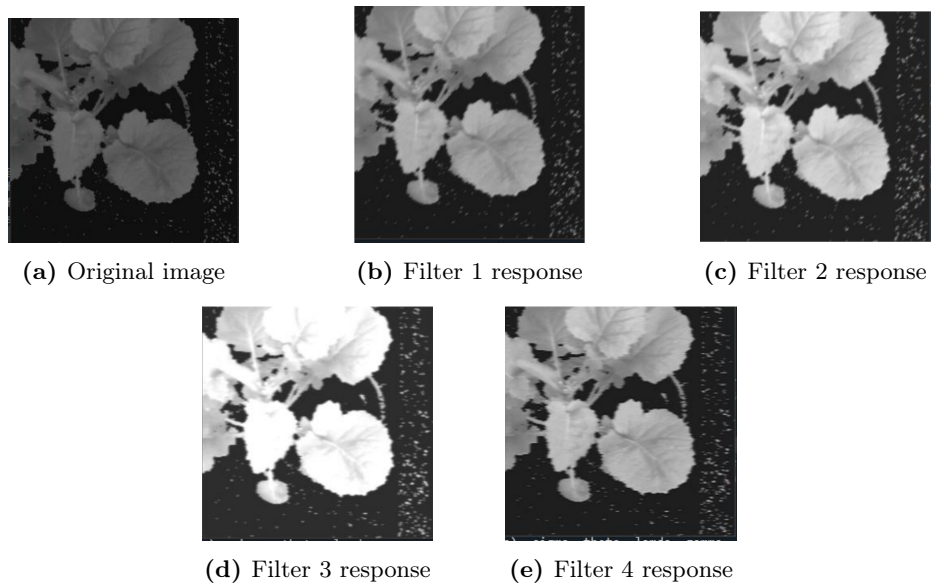


Fig. 3. Feature responses from the Gabor filter bank applied to a plant image.

training and 10% for testing. This resulted in training sets of 4,831 images

(canola) and 4,094 images (radish), with corresponding test sets of 547 and 525 images, respectively.

During hyperparameter optimisation, we further split the training data using stratified sampling, reserving 15% as a validation set for model selection while using the remaining 85% for model training. Feature extraction was performed on all images, but hyperparameter tuning of the classification methods (RBF-SVM and XGBoost) was performed using 3-fold stratified cross-validation in the training set to prevent data leakage. The validation set guided hyperparameter selection, while the test set was held out and untouched during model development for final performance evaluation. After selecting the optimal hyperparameters, each model was retrained on the entire training set and then evaluated once on the held-out test set.

The image input sizes were adjusted to fit the dimensionality requirements of the feature extractor methods proposed in this work. The input dimensionality of the CNN models was set to $224 \times 224 \times 1$, while that of the Gabor method was set to $128 \times 128 \times 1$. After preprocessing, the dataset sizes per class are summarised in Table 2.

Table 2. The number of images in the two species datasets at different growth stages.

| Class | Canola dataset | Radish dataset |
|----------------|----------------|----------------|
| Background | 2387 | 799 |
| Growth stage 1 | 527 | 782 |
| Growth stage 2 | 360 | 780 |
| Growth stage 3 | 1629 | 822 |
| Growth stage 4 | 475 | 1436 |
| Total | 5378 | 4619 |

4.2 Dataset visualisation

We used principal component analysis (PCA) [1] and t-SNE [14] to visualise the structure and correlation between each of the growth stages per plant class in the dataset. t-SNE reduces data dimensionality such that data points closer together in the original high-dimensional space remain close in the two-dimensional representation. The dimensionality reduction process is entirely unsupervised; labels are only used to colour-code the nodes when plotting. PCA was used as a preprocessing step to reduce dimensions before applying t-SNE to visualise the growth stage clusters. A t-SNE visualisation indicates whether classes can be easily separated without the use of feature extraction methods.

Figure 4a and Figure 4b show the two-dimensional visualisation of the canola dataset and the radish dataset created from the “bccr-segset” dataset. The visualisations show that the background and growth stage 4 classes can be distin-

gushed more clearly, but that the other classes, stages 1, 2, and 3, have overlapping patterns. Classifying these classes without identifying and extracting more nuanced features is expected to produce high misclassification rates.

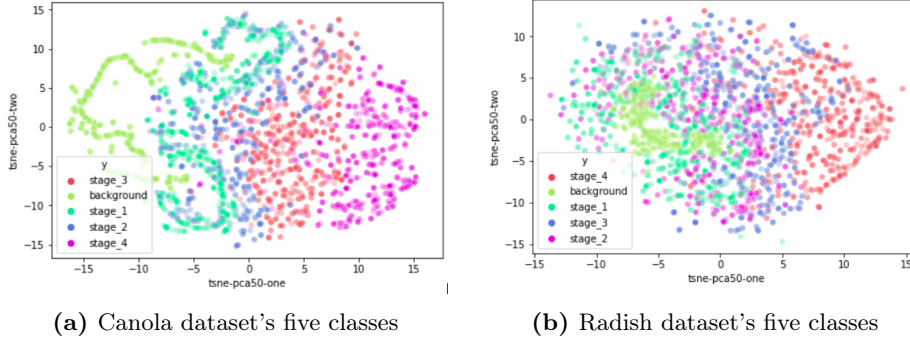


Fig. 4. Visualisation of the five classes for both datasets.

5 Classification

We first describe the process to select classification methods (Section 5.1), before discussing the implementation and optimisation of each method (Section 5.2).

5.1 Method Selection

After extracting high-level features via CNN or the Gabor pipeline, we selected SVMs and XGBoost over a standard multilayer perceptron (MLP) for classification due to the following technical considerations:

Regularisation and Limited Data: Our datasets were moderate in size, so overfitting was a concern. Fully connected neural layers (MLPs) introduce a large number of trainable parameters and tend to overfit when training data is very limited or noisy [23]. In contrast, SVMs and XGBoost have strong built-in regularisation (maximum-margin classification for SVM, and shrinkage/ensemble averaging for XGBoost).

Inference Speed and Efficiency: The SVM and XGBoost classifier ensures quick inference, since SVMs require only support vectors and kernel parameters, while XGBoost uses tree-based structures with inherently fewer parameters than dense neural network layers. SVMs compute a dot product with support vectors, and XGBoost uses simple tree traversals; both operations are efficient. Wang et al. [23] demonstrated that using an SVM as the classification layer reduced inference latency by 31% compared to a CNN with a softmax layer.

We chose SVM and XGBoost over other classical approaches, as SVMs are effective in high-dimensional feature spaces, such as those produced by CNN embeddings or Gabor filters, and incorporate explicit regularisation that improves generalisation [6]. XGBoost was preferred because gradient boosting reduces bias by focusing on difficult-to-classify samples, often leading to improved accuracy and compact models while maintaining efficient inference [3].

5.2 Implementation and Hyperparameter Optimisation

RBF-SVM Implementation RBF-SVM maps non-linear input data into a higher-dimensional space where linear separation is possible. Two hyperparameters control the RBF-SVM’s behaviour: the regularisation parameter C (penalty for misclassification) and γ (kernel spread). Higher C values yield smaller margins, while larger γ values create more flexible decision boundaries.

We optimised the RBF-SVM hyperparameters using scikit-learn’s RandomizedSearchCV with 3-fold stratified cross-validation on a representative training subset. The search ranges were $C \in \{2^{-2}, \dots, 2^6\}$ and $\gamma \in \{2^{-10}, \dots, 2^{-2}\}$. After 30 iterations, the optimal values were $C = 32$ and $\gamma = 0.00001$.

XGBoost Implementation XGBoost employs ensemble learning by combining multiple decision trees, enhanced with L1 and L2 regularisation to improve generalisation. The algorithm handles missing values naturally and often performs well on structured data, characteristics that seemed appealing to our feature-based approach. Following the same optimisation procedure as before, we tuned five key hyperparameters. Table 3 presents both search spaces and optimal values discovered through RandomizedSearchCV.

Table 3. XGBoost hyperparameter search space and optimal values.

| Hyperparameter | Search Space | Optimum Value |
|------------------|--------------|---------------|
| Learning rate | [0.2 to 0.7] | 0.5 |
| Max depth | [5 to 12] | 8 |
| Min child weight | [1 to 7] | 7 |
| gamma | [0.3 to 0.7] | 0.6 |
| Colsample bytree | [0.3 to 0.7] | 0.7 |

6 Results and Discussion

This section discusses the evaluation process and compares the results achieved by the two methods. We begin by visualising the extracted features using t-SNE and PCA for qualitative analysis, then conduct a quantitative comparison using accuracy, confusion matrices, and execution time as metrics.

6.1 Feature Visualisation and Qualitative Analysis

Figure 5 and Figure 6 show t-SNE-PCA visualisations of the different feature extraction approaches (Gabor filters, VGG-16, VGG-19, and ResNet-50) on both datasets, with each colour representing a different growth stage or background class.

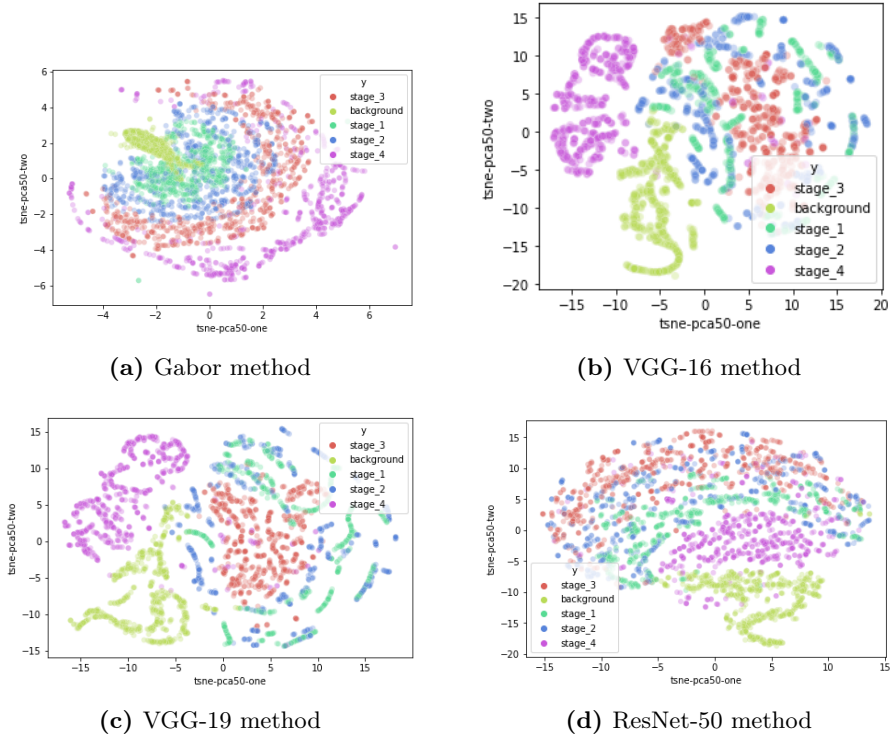


Fig. 5. Feature distributions of the canola training dataset using Gabor, VGG-16, VGG-19, and ResNet-50 feature extraction methods.

It is observed that VGG-16 and VGG-19 extract features that can be projected into separable spaces, while the Gabor method features also demonstrate reasonable separability. However, ResNet-50 showed lower clustering quality with more overlap between growth stages, likely because its residual connections, while effective for very deep networks, preserve low-level features that may act as noise rather than discriminative features for subtle plant texture variations. Despite this, all feature extraction methods appear capable of capturing features useful for growth stage classification, though some overlap remains visible across all techniques. This suggests that non-linear classifiers are necessary for effective classification.

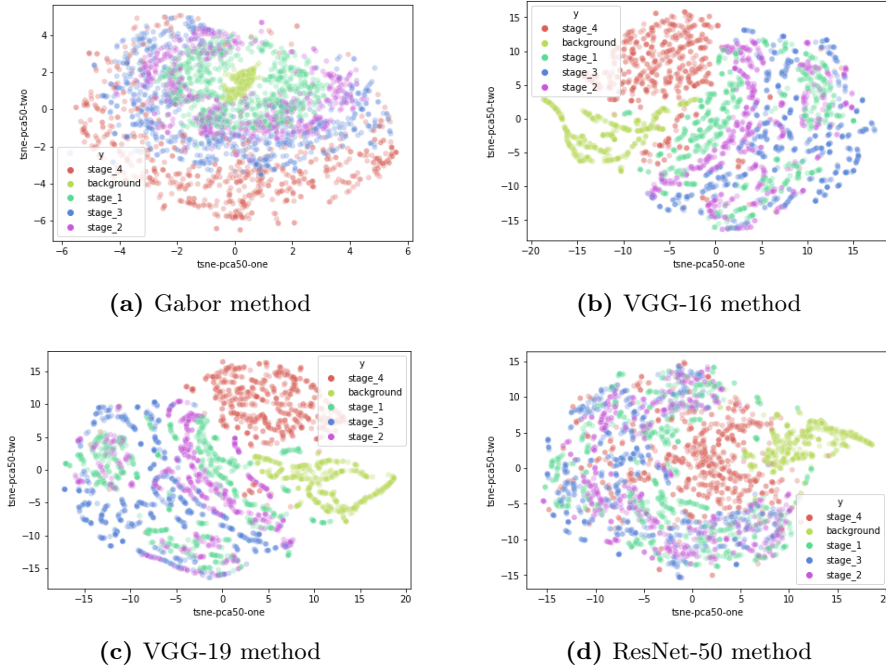


Fig. 6. Feature distributions of the radish training dataset using Gabor, VGG-16, VGG-19, and ResNet-50 feature extraction methods.

6.2 Classification Accuracy Evaluation

We evaluated the classification accuracy using three-fold stratified cross-validation. Figure 7 shows boxplots of classification accuracy on both test datasets for various model combinations. VGG-19 concatenated with RBF-SVM achieved the highest accuracy of 98.4% on both test datasets, demonstrating consistent performance across plant species. The Gabor filter method obtained competitive accuracies of 96.7% on canola and 91.4% on radish. ResNet-50 consistently achieved the lowest classification accuracy with both classification algorithms. The box plots also reveal the stability of the performance: VGG-19 showed consistent results across the folds, while other methods exhibited some variance.

Confusion matrices provide additional insight into classification performance. Figure 8 and Figure 9 summarise the results for both datasets. A notable pattern emerges: most models struggle with adjacent growth stages, which arise from gradual morphological transitions. For canola, misclassification occurs primarily between stages 1 and 2, where early stages share similar cotyledon structures and discriminative textures emerge only as true leaves develop. For radish, false negatives and false positives are frequent between stages 2 and 3. Transfer learning methods generally showed fewer classification errors compared to the Gabor

Feature extraction for plant growth estimation

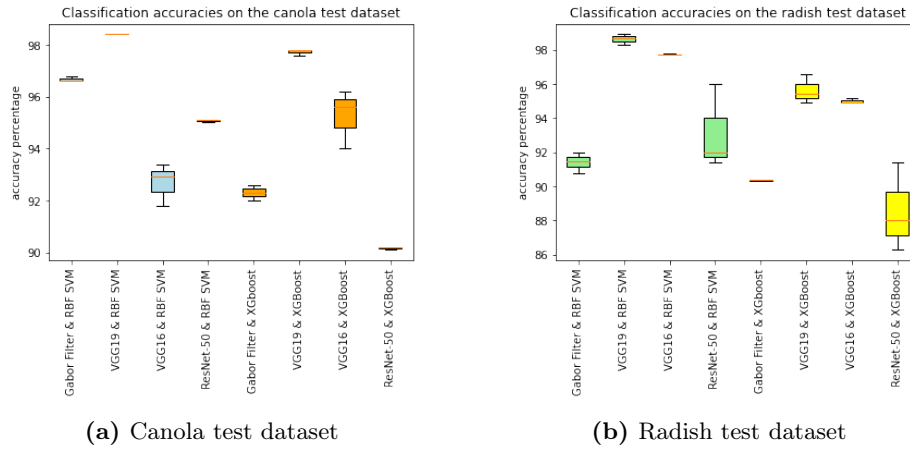


Fig. 7. Classification accuracies on the test datasets.

method, indicating that this represents a fundamental visual limitation rather than a methodological weakness.

6.3 Execution Time Analysis

For real-time prediction, processing time is important. We tested execution time for all feature extraction methods using an HP ProBook 450 7th Generation laptop. Table 4 shows the average testing time per image for both datasets.

The Gabor filter-based models required more time per image than CNN-based models, with the best case (Gabor filter & RBF-SVM) averaging 0.11-0.12 seconds per image. The most accurate model (VGG-19 & RBF-SVM) achieves 0.08 seconds per image on both datasets. These prediction times were obtained without GPU acceleration, suggesting that both approaches are suitable for real-time implementation. The Gabor filter-based method takes longer due to the

Table 4. Testing times of all models on both datasets (seconds per image).

| Method | Canola | Radish |
|------------------------|--------|--------|
| VGG-16 & RBF SVM | 0.07 | 0.06 |
| VGG-16 & XGBoost | 0.08 | 0.07 |
| VGG-19 & RBF SVM | 0.08 | 0.08 |
| VGG-19 & XGBoost | 0.09 | 0.09 |
| ResNet-50 & RBF SVM | 0.09 | 0.09 |
| ResNet-50 & XGBoost | 0.10 | 0.11 |
| Gabor Filter & RBF SVM | 0.12 | 0.11 |
| Gabor Filter & XGBoost | 0.30 | 0.31 |

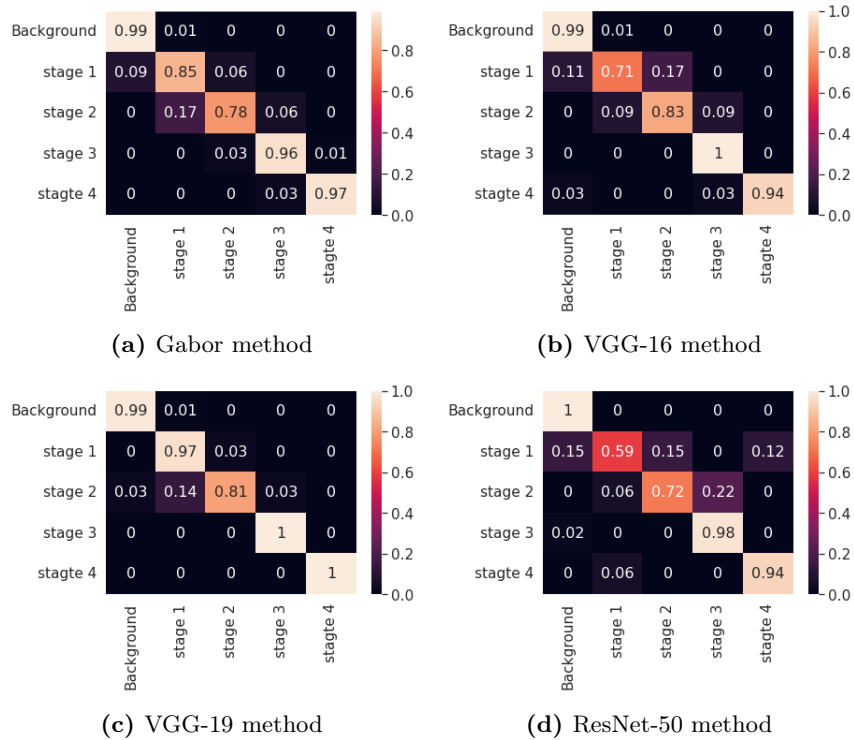


Fig. 8. Confusion matrices for the canola test dataset. Rows represent true classes, columns represent predictions.

morphological operations required before feature extraction, while CNN-based transfer learning models achieve faster processing since the convolutional layers used as feature extractors have been pre-trained.

7 Conclusion

In this study, we applied two feature extraction methods to estimate plant growth stages: one that uses a bank of Gabor filters and morphological operators, and another based on transfer learning from pre-trained CNNs. We evaluated their performance experimentally using two custom datasets of two plant species derived from the publicly available "bccr-segset" dataset. We fine-tuned two classification algorithms, RBF-SVM and XGBoost, and concatenated them with feature extraction methods to create a complete growth stage classification pipeline.

We used PCA and t-SNE to visualise growth stage image separation in data sets, before and after feature extraction. Initially, clustering showed significant overlap among growth stages, highlighting similar plant features across classes and the need for effective feature extraction. The improved clustering

Feature extraction for plant growth estimation

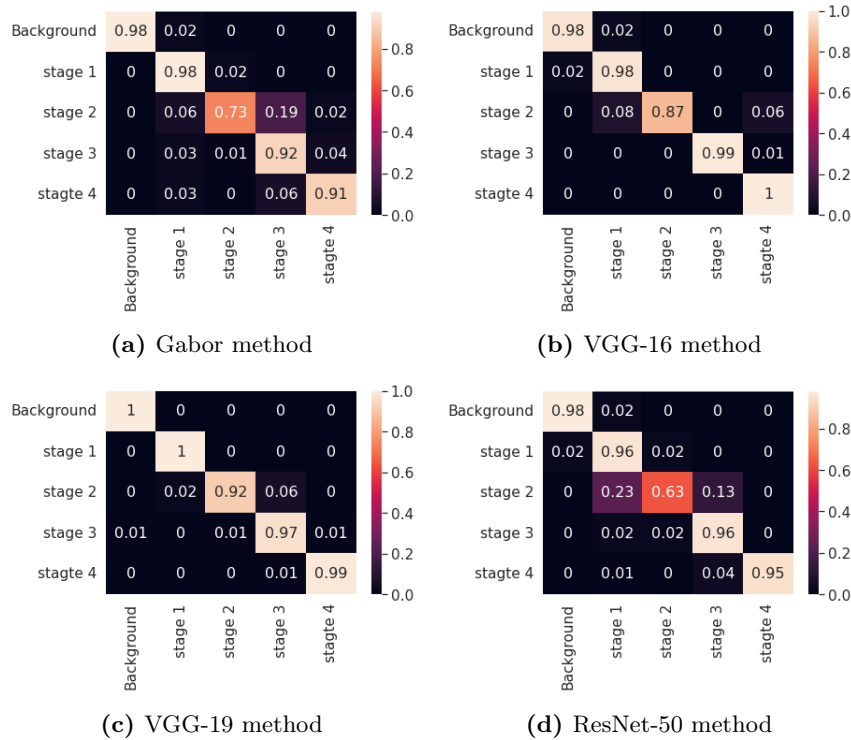


Fig. 9. Confusion matrices for the radish test dataset. Rows represent true classes, columns represent predictions.

after applying our feature extraction methods demonstrates their capability to extract features that enable accurate classification with minimal visible overlap between classes. Experimental results showed that the CNN-based transfer learning method can achieve an average accuracy of 98.4% in estimating four growth stages, while the Gabor method can achieve an average classification accuracy of 94.05%. Using the same hardware setup, we also tested the execution time of both approaches. The VGG-16 SVM feature extractor had the fastest execution time per image, followed by the VGG-19 feature extractor. The Gabor filter SVM method had a higher execution time than the CNN-based methods due to the additional morphological operators. For real-time applications, both approaches demonstrate practical feasibility.

Acknowledgements This work was partially supported by the Telkom CoE at NWU. In addition, this work is based on the research supported in part by the National Research Foundation of South Africa (Grant Reference Numbers: RCDL240215206999).

References

1. Abdi, H., Williams, L.J.: Principal Component Analysis. *Wiley Interdisciplinary Reviews: Computational Statistics* **2**(4), 433–459 (2010)
2. Alejandrino, J., Concepcion, R., Lauguico, S., Tobias, R.R., Almero, V.J., Puno, J.C., Bandala, A., Dadios, E., Flores, R.: Visual Classification of Lettuce Growth Stage Based on Morphological Attributes Using Unsupervised Machine Learning Models. In: 2020 IEEE Region 10 Conference (TENCON). pp. 438–443. IEEE (2020)
3. Chen, T., Guestrin, C.: XGBoost: A Scalable Tree Boosting System. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 785–794. ACM (2016)
4. Concepcion II, R.S., Alejandrino, J.D., Lauguico, S.C., Tobias, R.R., Sybingco, E., Dadios, E.P., Bandala, A.A.: Lettuce Growth Stage Identification Based on Phytomorphological Variations Using Coupled Color Superpixels and Multifold Watershed Transformation. *International Journal of Advances in Intelligent Informatics* **6**(3), 261–277 (2020)
5. Cope, J.S., Remagnino, P., Barman, S., Wilkin, P.: Plant Texture Classification Using Gabor Co-Occurrences. In: International Symposium on Visual Computing. pp. 669–677. Springer (2010)
6. Cortes, C., Vapnik, V.: Support-Vector Networks. *Machine Learning* **20**(3), 273–297 (1995)
7. Daugman, J.G.: Image Analysis by Local 2-D Spectral Signatures. *Journal of the Optical Society of America A* **2**, 74 (1985)
8. Gabor, D.: Theory of Communication. Part 1: The Analysis of Information. *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering* **93**(26), 429–441 (1946)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
10. Ji, Z., Pan, Y., Zhu, X., Wang, J., Li, Q.: Prediction of Crop Yield Using Phenological Information Extracted from Remote Sensing Vegetation Index. *Sensors* **21**(4) (2021). <https://doi.org/10.3390/s21041406>, <https://www.mdpi.com/1424-8220/21/4/1406>
11. Le, V.N.T., Ahderom, S., Alameh, K.: Performances of the LBP Based Algorithm over CNN Models for Detecting Crops and Weeds with Similar Morphologies. *Sensors* **20**(8), 2193 (2020)
12. Le, V.N.T., Apopei, B., Alameh, K.: Effective Plant Discrimination Based on the Combination of Local Binary Pattern Operators and Multiclass Support Vector Machine Methods. *Information Processing in Agriculture* **6**(1), 116–131 (2019)
13. Logeshwaran, J., Srivastava, D., Kumar, K.S., Rex, M.J., Al-Rasheed, A., Getahun, M., Soufiene, B.O.: Improving Crop Production Using an Agro-Deep Learning Framework in Precision Agriculture. *BMC Bioinformatics* **25**(1), 341 (2024)
14. van der Maaten, L., Hinton, G.: Visualizing Data Using t-SNE. *Journal of Machine Learning Research* **9**, 2579–2605 (2008)
15. Meraj, T., Sharif, M.I., Raza, M., Alabrah, A., Kadry, S., Gandomi, A.H.: Computer Vision-Based Plants Phenotyping: A Comprehensive Survey. *Iscience* **27**(1) (2024)
16. Mercier, F., Couasnet, G., El Ghaziri, A., Bouhleb, N., Sarniguet, A., Marchi, M., Barret, M., Rousseau, D.: Deep-Learning-Ready RGB-Depth Images of Seedling Development. *Plant Methods* **21**(1), 16 (2025)

17. Mukti, I.Z., Biswas, D.: Transfer Learning Based Plant Diseases Detection Using ResNet50. In: 2019 4th International Conference on Electrical Information and Communication Technology (EICT). pp. 1–6. IEEE (2019)
18. Murata, K., Ito, A., Takahashi, Y., Hatano, H.: A Study on Growth Stage Classification of Paddy Rice by CNN Using NDVI Images. In: 2019 Cybersecurity and Cyberforensics Conference (CCC). pp. 85–90. IEEE (2019)
19. Rasti, S., Bleakley, C.J., Silvestre, G., Holden, N., Langton, D., O'Hare, G.M.: Crop Growth Stage Estimation Prior to Canopy Closure Using Deep Learning Algorithms. *Neural Computing and Applications* **33**(5), 1733–1743 (2021)
20. Rezende, E., Ruppert, G., Carvalho, T., Theophilo, A., Ramos, F., Geus, P.d.: Malicious Software Classification Using VGG16 Deep Neural Network's Bottleneck Features. In: *Information Technology - New Generations*, pp. 51–59. Springer (2018)
21. Rimiru, R.M., Gateri, J., Kimwele, M.W.: GaborNet: Investigating the Importance of Color Space, Scale and Orientation for Image Classification. *PeerJ Computer Science* **8**, e890 (2022)
22. Salman, M., Mathavan, S., Kamal, K., Rahman, M.: Pavement Crack Detection Using the Gabor Filter. In: 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013). pp. 2039–2044. IEEE (2013)
23. Sima, R., Zhu, B., Wang, F., Wang, Y., Zhang, Z., Li, C., Wu, Z., Hong, B.: Real-Time Classification of Distributed Fiber Optic Monitoring Signals Using a 1D-CNN-SVM Framework for Pipeline Safety. *Processes* **13**(6) (2025). <https://doi.org/10.3390/pr13061825>, <https://www.mdpi.com/2227-9717/13/6/1825>
24. Simonyan, K., Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv preprint arXiv:1409.1556 (2014)
25. Soussi, A., Zero, E., Sacile, R., Trincherro, D., Fossa, M.: Smart Sensors and Smart Data for Precision Agriculture: A Review. *Sensors* **24**(8) (2024), <https://www.mdpi.com/1424-8220/24/8/2647>
26. Van Delden, S., SharathKumar, M., Butturini, M., Graamans, L., Heuvelink, E., Kacira, M., Kaiser, E., Klamer, R., Klerkx, L., Kootstra, G., et al.: Current status and future challenges in implementing and upscaling vertical farming systems. *Nature Food* **2**(12), 944–956 (2021)
27. Vyas, R., Kanumuri, T., Sheoran, G.: Iris Recognition Using 2-D Gabor Filter and XOR-SUM Code. In: 2016 1st India International Conference on Information Processing (IICIP). pp. 1–5. IEEE (2016)
28. Wäldchen, J., Mäder, P.: Plant Species Identification Using Computer Vision Techniques: A Systematic Literature Review. *Archives of Computational Methods in Engineering* **25**(2), 507–543 (2018)
29. Waqas, M., Naseem, A., Humphries, U.W., Hlaing, P.T., Dechpichai, P., Wangwongchai, A.: Applications of Machine Learning and Deep Learning in Agriculture: A Comprehensive Review. *Green Technologies and Sustainability* p. 100199 (2025)