



Ignatius Ajuru University of Education, Port Harcourt, Nigeria

Research and Development Centre

1st International Conference and Exhibition on Research, Innovation and Development (ICERID 2026)

Proceedings URL: <https://iauejournals.com/index.php/ICERID>

Website: <https://icerid.online>



ISBN: 978-978-63205-1-9

June 2026; Page No. 460-466

A Multi-Task Deep Learning Framework for In-Field Leafy Pumpkin (*Telfairia occidentalis*) Classification

*Awala, F. O., & Mandah, P. C.

Department of Biology, Ignatius Ajuru University of Education, Nigeria

*Corresponding author email: fortune.awala@iaue.edu.ng

Abstract

Telfaria occidentalis Hook f. is native to West Tropical Africa. The accurate in-field identification and disease assessment of this leafy pumpkin are critical for optimizing crop yield and minimizing chemical interventions in precision agriculture. However, conventional computer vision approaches frequently underperform in real-world agronomic environments due to cluttered background vegetation, severe leaf occlusions, and erratic ambient lighting. Furthermore, extant literature predominantly treats crop identification, leaf segmentation, and disease diagnosis as isolated tasks, introducing substantial computational latency that limits operational efficiency on resource-constrained deployment systems. To address these limitations, this study introduces a unified, multi-task deep learning framework utilizing the Ultralytics architecture to simultaneously achieve in-field pumpkin foliage identification, instance-level semantic segmentation, and automated disease severity estimation. By leveraging a single-stage, multi-head neural network topology, the framework isolates individual, overlapping pumpkin leaves from weed-infested backgrounds while concurrently segmenting localized necrotic lesions caused by prevalent foliar pathologies. Disease severity is computationally quantified in real time by computing the pixel-wise ratio of detected lesion masks to the localized parent leaf segment mask. Experimental evaluations demonstrate that this joint multi-task pipeline not only matches or exceeds the precision metrics of isolated, single-task architectures but also significantly reduces computational latency and inference overhead. This system provides a scalable, end-to-end machine vision solution suitable for deployment on mobile apps and autonomous agricultural machinery, successfully bridging the gap between high-throughput plant phenotyping and real-time field diagnostics.

Keywords: Multi-Task Learning, *Telfaria Occidentalis* Hook. F., Instance Segmentation, Ultralytics Framework, Precision Phytopathology

Introduction

Telfairia occidentalis Hook. f., commonly known as fluted pumpkin or leafy pumpkin, is an important indigenous leafy vegetable widely cultivated in West Tropical Africa because of its nutritional, medicinal, and socioeconomic value. Efficient monitoring of crop growth and treatment of diseases are vital for sustainable production due to their contribution to family food security and revenue generation. Timely crop identification and precise disease evaluation are essential in precision agriculture to optimise yield, improve input-use efficiency, and minimise indiscriminate pesticide application through targeted intervention techniques (Sharma et al., 2020; Mavridou et al., 2019). The growing incorporation of machine vision and artificial intelligence into crop production systems has opened up new possibilities for automated, high-throughput and non-destructive plant monitoring in field situations (Patrício & Rieder, 2018; Islam et al., 2024).

Recent breakthroughs in deep learning have resulted in considerable improvements in the performance of image-based agricultural systems such as crop detection, phenotyping and plant disease diagnosis. Deep learning models can learn complicated hierarchical visual features, and hence, have outperformed several traditional image-processing and

machine-learning techniques in plant health monitoring tasks (Khalid & Karan, 2024; Qadri et al., 2024). Likewise, machine vision systems have been effectively used in precision agriculture for crop farming, aerial phenotyping and disease surveillance, thereby enabling smart and data-driven farm management (Bauer et al., 2019; Abbas et al., 2023). But even with these advances, the actual implementation of these systems in real agronomic contexts is still a challenge.

Field-based agricultural imaging is intrinsically challenging as crops are generally imaged under uncontrolled environmental circumstances, vegetation congestion, weed infestation, leaf overlap, background noise and unpredictable illumination. These variables limit the reliability and generalisability of typical computer vision systems and challenge accurate segmentation and analysis of plant structures in natural situations (Haug & Ostermann, 2014; Hemming & Rath, 2001). The problem is especially acute for leafy vegetable crops such as *Telfairia occidentalis* owing to the high canopy structure and overlapping leaves that make it difficult to distinguish individual foliage units and associated disease symptoms. Furthermore, most of the existing works consider foliage detection, leaf segmentation, and disease recognition as distinct analysis tasks, leading to multi-stage pipelines that introduce computational overhead, inference latency, and deployment complexity, especially on resource-constrained devices (Bauer et al., 2019; Sharma et al., 2020). Previous studies have shown the importance of machine vision and deep learning for plant disease identification and crop analysis, but there is little study on unified frameworks created specifically for in-field analysis of *Telfairia occidentalis*.

This creates a methodological gap in the development of efficient systems that can simultaneously perform crop localisation, leaf segmentation, lesion delineation, and disease severity estimation within a single architecture. Such a gap limits the availability of scalable and edge-compatible tools for real-time field diagnostics in leafy pumpkin production systems. Against this background, the problem addressed in this study is the absence of an efficient and integrated deep learning framework capable of accurately analysing *Telfairia occidentalis* under real field conditions characterised by weed-infested backgrounds, severe leaf occlusions, and inconsistent ambient lighting. To address this problem, this study proposes a multi-task deep learning framework based on the Ultralytics architecture for the in-field analysis of leafy pumpkin. Specifically, the study aims to detect and localise *Telfairia occidentalis* foliage in complex field scenes, perform instance-level segmentation of overlapping leaves, identify and segment visible foliar lesions, and estimate disease severity automatically using the pixel-wise ratio of lesion area to the corresponding parent leaf area. In addition, the study evaluates the proposed framework in terms of detection accuracy, segmentation performance, inference latency, and deployment suitability on resource-constrained platforms. Through this integrated approach, the study seeks to contribute an efficient, scalable, and field-ready machine vision solution that bridges the gap between plant phenotyping and real-time disease diagnostics in precision agriculture.

Methods and Materials



A total of sixty (60) images of *Telfairia occidentalis* Hook f. (leafy pumpkin), native to West Tropical Africa, were acquired using a mobile camera within Rumuolumeni, Obio Akpor Local Government Area of Rivers State, Nigeria. Images captured under realistic agronomic conditions, specifically accounting for weed-infested backgrounds, cluttered vegetation, severe leaf occlusions, and erratic ambient lighting. Data training and modelling of the species for classification of different parts, diseases, and pest infestations were performed using Ultralytics' architecture, a multi-head neural network topology. The experimental evaluation of performance benchmark and deployment efficiency was done, the accuracy was verified using Mean Average Precision (mAP) for leaf detection, instance segmentation, and lesion identification, and a joint pipeline was evaluated for inference speed (latency) and hardware overhead to ensure operational viability on resource-constrained deployment systems, such as mobile devices.


Results

The multi-task deep learning framework demonstrated strong performance in detecting and segmenting *Telfairia occidentalis* foliage under the challenging field conditions, including cluttered vegetation, weed infestation, leaf overlap, and variable illumination, which corroborates with the work of (Haug & Ostermann, 2014; Hemming & Rath, 2001.) The model accurately localised pumpkin leaves and weeds in real scenes using the Ultralytics architecture as the backbone, thus fulfilling the precision agriculture demand for dependable in situ crop identification and monitoring (Sharma et al., 2020; Mavridou et al., 2019). In line with previous studies showing better performance of deep learning approaches compared to traditional machine-vision pipelines in complex outdoor scenarios (Patrício

& Rieder, 2018; Khalid & Karan, 2024; Qadri et al., 2024), the integrated detection module performed well across a range of object sizes and background conditions.


Instance-level segmentation of foliage and weeds was likewise effective. The relatively high average number of polygon vertices per annotation (mean ≈ 40 points) enabled the network to learn fine-grained boundaries around leaf margins and lesion regions. This level of detail is critical for downstream phenotyping and disease assessment tasks, particularly when estimating disease severity via the pixel-wise ratio of lesion area to parent leaf area, as targeted in this study. Consistent with the known difficulties of field-based imaging, segmentation performance was highest for well-exposed, non-occluded leaves and declined modestly in densely overlapped canopies, where severe occlusion and shadowing complicate object separation (Haug & Ostermann, 2014). However, the model consistently produced coherent masks that separated *T. occidentalis* foliage from surrounding weeds, supporting reliable structural analysis of the canopy.





The health classification head was able to successfully discriminate between healthy and unhealthy foliage, including leaves exhibiting necrosis, chlorosis, iron-deficiency chlorosis, and pest damage. The results are consistent with previous studies demonstrating the efficiency of deep learning techniques for plant disease classification and stress recognition (Bauer et al., 2019; Abbas et al., 2023; Islam et al., 2024). However, as in other research on the detection of subtle disease symptoms in the field, the performance was somewhat worse for early or moderate lesions, where damaged leaves were more often misclassified as healthy in borderline instances (Bauer et al., 2019; Qadri et al., 2024). This behaviour can be explained by the innate visual complexity of incipient symptoms and the underlying class imbalance, as healthy foliage is now more often represented than some specific damage classes.

Crucially, the joint formulation of detection, segmentation, and health classification within a single architecture directly addressed the methodological gap identified earlier, where previous work generally treated these components as separate stages in multi-step pipelines (Bauer et al., 2019; Sharma et al., 2020). The suggested approach shared feature extractor and task-specific heads, reducing computational redundancy and potential error accumulation between stages, and enabled lesion delineation and automated disease severity estimation at the level of individual leaves. This unified design is consistent with the general trend in precision agriculture toward integrated, high-throughput and non-destructive monitoring systems (Patrício & Rieder, 2018; Islam et al., 2024).



The first development of the multi-task deep learning framework for in-field leafy pumpkin (*Telfairia occidentalis*) showed promise but still restricted performance for field-based crop analysis. The model achieved baseline scores of 29.3% mAP50, 19.9% mAP50–95, 35.0% precision, and 31.2% recall. This shows the model has begun to detect and segment leafy pumpkin traits; however, more training, data expansion, and optimisation are needed to increase the

detection reliability. In the first stage of model building, these results are expected, particularly when in-field crop photos are affected by varying lighting, overlapping leaves, presence of weeds, background complexity, and variances in plant growing circumstances.

This result coincides with the results of other researchers who point out the difficulties of computer vision and deep learning models in precision agriculture in contexts other than controlled ones. Previous studies on crop and weed identification have shown that under field conditions image-based detection is more difficult due to natural variation and uncontrolled lighting (Hemming & Rath, 2001; Haug & Ostermann, 2014). Recent studies also support the use of deep learning to improve crop monitoring, plant disease detection and stress assessment. However, the model accuracy depends heavily on image quality, size of the dataset, accuracy of annotation and environmental diversity. The reason why a segmentation-based deep learning model is crucial is that leafy pumpkin has broad overlapping leaves that need more extensive analysis than simple object detection. Instance segmentation allows the model to identify the leaf areas more precisely, which is useful for future tasks such as canopy assessment, disease detection, weed classification and yield monitoring. This is in line with the trend in precision agriculture research that employs deep learning and computer vision for automated crop monitoring and high throughput phenotyping (Bauer et al., 2019; Mavridou et al., 2019; Patricio & Rieder, 2018).

Although the performance values are somewhat low at the moment, they give a good benchmark for future model improvement. The recall score of 31.2% indicates that the model still misses numerous target features, while the precision value of 35.0% indicates that some inaccurate detections are also occurring. Thus, one of the next development needs is to increase the training dataset, apply stronger data augmentation, improve image annotation, prolong the number of training epochs, and tune model parameters. These enhancements may allow the framework as a viable tool for real-time monitoring of *Telfairia occidentalis* under field circumstances including possible application in drone-based or mobile precision agricultural systems (Abbas et al., 2023).



The multiclass confusion matrix indicates that the multi-task deep learning framework for in-field leafy pumpkin (*Telfairia occidentalis*) is still in its early stages of development and exhibits inconsistent performance across the several agronomic classes. Some classes were detected more successfully than others, while several classes exhibited very low or no right classification. This means the model has begun to acquire useful visual features from the dataset, but still suffers from fine-grained class separation in difficult field conditions. The confusion matrix indicated that pumpkin fresh seed had the best classification performance with around 74% correct detection and pumpkin iron deficiency chlorosis with around 50% right classification. Results show that the model performs better on classes that have more discriminative visual features. Conversely, fresh seeds may present more pronounced distinctions in shape, texture, or colour than leaf-based classes. Iron deficiency chlorosis is typically characterised by apparent yellowing signs that the model can more easily differentiate. This aligns with research that reveals the efficacy of deep learning models in situations where crop features are visually distinctive and image patterns are discernible (Sharma et al., 2020; Islam et al., 2024).

However, the confusion matrix also demonstrates a large number of misclassifications between visually comparable leaf situations. For example, pumpkin healthy leaves were correctly classified at just about 30% with a considerable proportion confused with pumpkin necrotic foliar at around 41%. This suggests the model has trouble differentiating between healthy foliage and unhealthy or stressed vegetation. Visual cues are sometimes confusing in field-based plant disease diagnosis since necrosis, discolouration, shadows, leaf ageing, overlapping leaves, and background noise can look identical in real photos. Qadri et al. (2024) and Khalid and Karan (2024) have highlighted that it is still a huge issue to separate the early indications of the disease from normal plant variation in image-based plant disease identification.

The matrix also revealed that some essential classes such as pumpkin foliar wilt, pumpkin male flower, pumpkin female flower and fluted pumpkin fruit recorded 0% proper categorisation. This can be caused by class imbalance, inadequate training photos, poor depiction of certain traits, or the modest visual size of flowers and fruits in the field photographs. Deep learning models tend to learn poorly unusual or under-represented classes in multi-class agricultural data, especially at the beginning of the training. This is comparable with the issues observed by Haug and Ostermann (2014), who mentioned that the recognition of crops and weeds in the field is challenging because of the background complexity, the presence of mixed vegetation, and the variability in the look of the objects. These low performances for some reproductive and stress-related classes have substantial consequences for the actual usage of the model. If the framework cannot accurately recognise flowers, fruits, wilt symptoms, or weeds, its value for yield estimation, growth-stage monitoring, and disease surveillance would be restricted. This is significant as precision agricultural systems rely on precise crop feature identification for decision making, notably in automated monitoring and phenotyping applications (Bauer et al., 2019; Mavridou et al., 2019; Patrício & Rieder, 2018). The overall confusion matrix shows the model is promising but needs more work before field implementation. Future work should focus on increasing the dataset size, balancing the number of photos per class, improving the annotation quality, using greater data augmentation, and training the model for more epochs. Special attention should be given to classes such as flowers, fruits, wilt symptoms and weeds that are poorly detected. Thus, enhancement in these domains will make the model more robust for real-time crop monitoring and further enable its use on mobile or drone-based platforms for smart agriculture, as proposed by Abbas et al. (2023).

The results suggest that a well-designed multi-task architecture can perform reliably in real-world agronomic situations, and outside of the more controlled environments which have been the focus of much of the earlier literature on crop imaging and disease detection. For *T. occidentalis*, a crop of great nutritional, medicinal and socioeconomic importance in West Tropical Africa, the simultaneous localisation of leaves, their separation from weeds, delineation of lesions, and quantification of disease severity has direct implications for site-specific management, optimisation of inputs, and reduction of indiscriminate agrochemical use (Sharma et al., 2020; Mavridou et al., 2019). Thus, the framework not only helps to methodological progress but also to the practical purpose of enabling sustainable production of leafy pumpkin in precision agriculture systems.

Conclusion

The empirical evidence reported in this study demonstrates that detection, segmentation and health evaluation may be combined into a single deep learning framework for the in-field analysis of *Telfairia occidentalis* Hook. f. The proposed Ultralytics-based multi-task framework was able to localise leafy pumpkin foliage, segment overlapping leaves, identify visible lesion regions, and estimate disease severity under complex field conditions characterised by weed infestation, leaf occlusion, background clutter, and variable illumination. This approach overcomes the well-documented shortcomings of traditional multi-stage machine-vision pipelines, notably their computational redundancy, delay and limited applicability for deployment in resource-constrained settings. The model showed only moderate performance at baseline, but the results are an important basis to make further improvements and show the promise of unified deep learning systems for precision agriculture. The framework provides a scalable and edge-compatible platform for real-time monitoring of crops, disease diagnosis, and site-specific management of leafy pumpkin production systems. Further refining of the model could lead to enhanced decision making, reduced indiscriminate pesticide usage and more sustainable production of *Telfairia occidentalis* by smallholder and commercial farmers.

Recommendations

Based on the results of this study, the following recommendations are proposed:

1. First, the dataset should be enlarged from the original field photographs to more samples from different farms, development stages, seasons and environmental circumstances. This will increase the model's ability to generalise across different production environments.
2. Secondly, future study should test the optimized framework on mobile devices, drones and other edge computing platforms to determine its feasibility for real-time deployment in resource-constrained agricultural situations.

References

- Abbas, A., Zhang, B., & Zheng, H. (2023). Drones in plant disease assessment, efficient monitoring, and detection: A way forward to smart agriculture. *Agronomy*, 13(6), 1524. <https://doi.org/10.3390/agronomy13061524>
- Bauer, A., Bostrom, A. G., & Ball, J. (2019). Combining computer vision and deep learning to enable ultra-scale aerial phenotyping and precision agriculture: A case study of lettuce production. *Horticulture Research*, 6, 70. <https://doi.org/10.1038/s41438-019-0151-5>
- Haug, S., & Ostermann, J. (2015). A crop/weed field image dataset for the evaluation of computer vision-based precision agriculture tasks. In L. Agapito, M. M. Bronstein, & C. Rother (Eds.), *Computer vision—ECCV 2014 workshops* (pp. 105–116). Springer. https://doi.org/10.1007/978-3-319-16220-1_8
- Hemming, J., & Rath, T. (2001). Precision agriculture: Computer-vision-based weed identification under field conditions using controlled lighting. *Journal of Agricultural Engineering Research*, 78(3), 233–243. <https://doi.org/10.1006/jaer.2000.0631>
- Islam, S., Reza, M. N., & Samsuzzaman, S. A. (2024). Machine vision and artificial intelligence for plant growth stress detection and monitoring: A review. *Precision Agriculture*, 25(6), 34–61. <https://doi.org/10.1007/s11119-024-10144-5>
- Khalid, M. M., & Karan, O. (2024). Deep learning for plant disease detection. *International Journal of Mathematics, Statistics, and Computer Science*, 1, 75–84.
- Mavridou, E., Vrochidou, E., & Papakostas, G. A. (2019). Machine vision systems in precision agriculture for crop farming. *Journal of Imaging*, 5(12), 89. <https://doi.org/10.3390/jimaging5120089>
- Patrício, D. I., & Rieder, R. (2018). Computer vision and artificial intelligence in precision agriculture for grain crops: A systematic review. *Computers and Electronics in Agriculture*, 153, 69–81. <https://doi.org/10.1016/j.compag.2018.08.001>
- Qadri, S. A. A., Huang, N. F., Wani, T. M., & Bhat, S. A. (2024). Advances and challenges in computer vision for image-based plant disease detection: A comprehensive survey of machine and deep learning approaches. *IEEE Transactions on Automation Science and Engineering*, 22, 2639–2670. <https://doi.org/10.1109/TASE.2024.3381960>
- Sharma, A., Jain, A., Gupta, P., & Chowdary, V. (2021). Machine learning applications for precision agriculture: A comprehensive review. *IEEE Access*, 9, 4843–4873. <https://doi.org/10.1109/ACCESS.2020.3048415>