

Morphological Database and Quantification of Liberica Coffee Beans Using a Hybrid CNN Model

Dioneces O. Alimoren¹, Francis G. Balazon¹, Jonnah R. Melo¹, Richelle Sulit¹ and Melmar Eje²

¹College of Informatics and Computing Sciences, Batangas State University, The National Engineering University Lipa Campus, Lipa City, Philippines

²College of Agriculture, Batangas State University, The National Engineering University Lobo Campus, Lobo, Batangas, Philippines

Article history

Received: 10-12-2025

Revised: 07-05-2026

Accepted: 15-07-2026

Corresponding Author:

Francis G. Balazon
Department of CICS, Batangas
State University the National
Engineering University Lipa
Campus, Philippines
Email: fbalazon@yahoo.com

Abstract: Traditional evaluation of Liberica coffee (Kapeng Barako) in Batangas is commonly performed through manual visual inspection and sensory judgment, which are slow, subjective, and difficult to reproduce across farms, buyers, and grading personnel. These limitations can lead to inconsistent quality classification, weak traceability, and reduced bargaining power for local producers. This study develops and validates a web-based morphological analysis system for Liberica coffee beans using a hybrid pipeline that combines U-Net convolutional neural network segmentation with deterministic feature extraction. A database of 4,000 high-resolution images collected from farms in Batangas was curated and divided into training, validation, and test subsets using a 70:15:15 ratio. To evaluate generalization, an independent 500-image dataset from previously unseen farms was also used. The model was trained with Dice plus binary cross-entropy loss, Adam optimization, and augmentation strategies designed to simulate field variations. The 96.0% accuracy reported in this study is explicitly defined as feature-level agreement: the percentage of test images in which the primary extracted traits fall within accepted tolerances relative to a manually verified reference. On the internal test set, the proposed method achieved a Dice coefficient of 0.971, intersection-over-union of 0.944, and high agreement for key morphological traits including area ($R^2 = 0.976$), eccentricity ($R^2 = 0.972$), solidity ($R^2 = 0.973$), and centroid coordinates ($R^2 > 0.962$). Compared with classical image processing and CNN baselines, the hybrid approach provided a favorable balance of accuracy, interpretability, and deployment speed. The resulting platform provides farmers, cooperatives, and researchers with a practical, low-cost, and reproducible tool for data-driven Liberica coffee quality assessment and morphology database development.

Keywords: Liberica Coffee, Kapeng Barako, U-Net, Convolutional Neural Network, Morphological Analysis, Computer Vision, Plant Phenotyping, Quality Assessment

Introduction

Liberica coffee (*Coffea liberica*), locally known as Kapeng Barako, is one of the coffee types grown in the Philippines and is associated with the rich, intense flavor profile described for Liberica or Barako coffee (Philippine Council for Agriculture, Aquatic and Natural Resources Research and Development [PCAARRD], 2023). In Batangas and neighboring coffee-producing areas, this crop has cultural and economic value because it supports local identity, farm livelihoods, specialty-market positioning, and cooperative-level product differentiation.

Traditional Liberica bean evaluation is usually based on manual inspection, visual sorting, and expert judgment. Although these practices remain useful, they are vulnerable to fatigue, lighting variability, evaluator

bias, and limited repeatability. Similar limitations have been reported in coffee inspection studies where manual grading is described as labor-intensive, error-prone, and difficult to scale (Garcia et al., 2019; Reales et al., 2024). For local producers, these limitations can result in inconsistent grading, weak traceability, reduced buyer confidence, and lower bargaining power.

Recent agricultural computer vision studies show that deep learning can support postharvest quality control, plant phenotyping, defect detection, and decision support (Murphy et al., 2024; Wang et al., 2025). In coffee, Convolutional Neural Network (CNN)-based models have been used for defect classification and coffee species or variety classification, including real-time or web-oriented applications (Arwatchananukul et al., 2024; Korkmaz et al., 2025; Laureta et al., 2025). However,

many systems focus on classification only, rely on opaque end-to-end predictions, or do not provide transparent morphological measurements that can be audited by farmers and researchers.

This study addresses that gap by integrating deep segmentation with deterministic morphometric computation. The CNN is used for the task it performs best - separating the bean from the background under varying image conditions - while final measurements are computed using explicit mathematical definitions. This design improves interpretability compared with direct feature regression and improves robustness compared with threshold-only processing. The web-based implementation translates the model into a practical tool for farmers, cooperatives, and researchers.

The core contributions of this work are:

- (1) A curated Liberica bean image database for quantitative morphology
- (2) A hybrid U-Net and deterministic feature extraction pipeline
- (3) An explicit definition of the reported 96.0% accuracy
- (4) Validation using train/validation/test partitioning, five-fold cross-validation, and an independent dataset
- (5) Baseline comparison against classical image processing and CNN-based alternatives
- (6) Deployment of the model in a web application that supports image upload, feature reporting, visualization, and downloadable records

Review of Related Literature

Deep learning has become a major tool for image-based plant phenotyping because it can handle complex object boundaries, variable backgrounds, and high-throughput image volumes. Murphy et al. (2024); Wang et al. (2025) emphasize that current phenotyping challenges include annotated-data scarcity, generalization under real-world imaging conditions, edge deployment, uncertainty estimation, and interpretable outputs. These issues motivate the present study because Liberica bean images must be analyzed under farm-relevant conditions and converted into measurements that are understandable to non-specialist users.

Coffee computer vision has progressed from classical machine vision and feature-based classifiers to CNNs and lightweight deployment models. Garcia et al. (2019) demonstrated that size, morphology, shape, and color information can support green coffee bean quality and defect inspection. More recent work by Arwatchananukul et al. (2024) used CNN architectures for Thai Arabica green coffee defect classification and demonstrated web-app deployment, while Korkmaz et al. (2025) compared multiple deep CNN architectures for coffee bean species classification. In the Philippine context, Reales et al.

(2024) presented a computer-vision approach for green coffee bean quality inspection, and Laureta et al. (2025) explored CNN-based coffee classification for Philippine coffee varieties.

Seed and bean morphology have long been quantified using image-analysis descriptors such as area, perimeter, eccentricity, solidity, roundness, and centroid location. SmartGrain showed that automated image analysis can calculate seed area, perimeter, and other shape traits at high throughput (Tanabata et al., 2012). The same principle supports coffee morphology quantification, but U-Net-style segmentation offers stronger boundary detection than thresholding alone. U-Net was originally developed for biomedical segmentation and is known for its encoder-decoder structure and skip connections that preserve localization detail (Ronneberger et al., 2015). Similar segmentation-based phenotyping tools have also been used in crop measurement systems such as strawberry phenotyping (Ndikumana et al., 2024).

Compared with a purely classical pipeline, the hybrid approach is less sensitive to shadows, slight background texture, and moderate illumination changes. Compared with a pure regression CNN, it avoids a black-box mapping from image to numeric features and allows each measurement to be visually checked through the segmentation mask. This combination is especially appropriate for agricultural deployment, where users need accuracy, speed, and trust in the basis of the reported measurements.

Methods

System Architecture and Workflow

The system was structured into a model-development phase and a deployment phase. During model development, curated bean images were preprocessed, augmented, used for U-Net training, and evaluated against verified reference measurements. During deployment, the trained model was embedded in a Flask-based web application in which a user uploads an image, the backend performs model inference, and the frontend displays segmentation and morphometric outputs. Fig. 1 is discussed here because it clarifies the complete path from image curation to real-time reporting, including how the offline training workflow differs from the online user-facing workflow.

Dataset Acquisition and Preprocessing

The primary dataset consisted of 4,000 high-resolution Liberica bean images collected from farms in Batangas Province. Images were acquired using a Samsung S24 main camera mounted on a tripod at a fixed distance from the sample. Beans were placed on a matte white background and illuminated using diffuse LED lighting to minimize shadows and specular reflections. All images were resized to 256x256 pixels and normalized to the [0, 1] pixel-intensity range.

TRAINING AND DEPLOYMENT WORKFLOW




Fig. 1: Model-development and deployment workflow for the Liberica bean morphology system

The dataset was split at the image level into 2,800 training images (70%), 600 validation images (15%), and 600 internal test images (15%). To reduce leakage, images from the same capture session were kept within the same split when possible. An additional independent validation set of 500 images was collected from farms not represented in training. This independent set was used only after model selection to evaluate generalization across new farm sources and slightly different image conditions.

Ground Truth Generation and Physical Verification

The initial reference masks and morphological measurements were generated using a classical image-processing pipeline consisting of grayscale conversion, Otsu thresholding, morphological opening/closing, contour filtering, and region-property extraction. To address the concern that the model would otherwise be validated only against another algorithm, the reference data were manually verified. A trained reviewer inspected the generated masks and corrected failures before inclusion in the reference CSV. In addition, a 200-bean subset was measured using digital calipers and manual ImageJ tracing to compare physical length-width estimates and image-derived area/perimeter values. The classical pipeline therefore, served as an annotation aid rather than an unverified ground-truth source.

The study acknowledges that pixel-based measurements are not a complete replacement for direct physical measurements unless spatial calibration is performed. For this reason, the current tool reports pixel units by default and uses a fixed-camera protocol. Future versions should incorporate a calibration marker in every

image to convert pixel measurements into millimeters with traceable physical units.

Data Augmentation

To improve robustness, online augmentation was applied only to the training set. The augmentation protocol included random rotation (± 25 degrees), horizontal and vertical flipping, width and height shifts up to 10%, zoom variation up to 15%, brightness variation of $\pm 20\%$, mild Gaussian noise, and small contrast changes. These operations simulate common image-acquisition differences while preserving the true morphological shape of the bean. The use of augmentation follows the general motivation of segmentation networks, which can benefit from data variation when annotated samples are limited (Ronneberger et al., 2015; Arwatchananukul et al., 2024).

Hybrid CNN Model and Feature Extraction

The segmentation model used a U-Net encoder-decoder architecture with skip connections. The encoder contained convolutional blocks with 32, 64, 128, and 256 filters, followed by a 512-filter bottleneck. Each block used 3×3 convolution, ReLU activation, batch normalization, and max pooling. The decoder used upsampling layers and skip connections to recover boundary detail. The output layer used a sigmoid activation to produce a 256×256 probability mask, which was binarized at a threshold of 0.5.

The hybrid nature of the method is central to the proposed contribution. The CNN estimates the bean mask, while deterministic scikit-image region-property functions compute area, perimeter, eccentricity, solidity, and centroid. The scikit-image library supports

reproducible image-analysis operations and region measurements in Python (Walt et al., 2014). Fig. 2 is discussed in this section because it shows both model architecture and the downstream deterministic extraction stage, making clear that the final features are not produced by an opaque regression head but by auditable mathematical measurements.

Training Protocol and Hyperparameters

The model was implemented in TensorFlow an open-source machine learning system for heterogeneous platforms (Abadi et al., 2016). It was trained using the Adam optimizer with an initial learning rate of 0.001 and ReduceLROnPlateau scheduling down to 0.00001. Adam was selected because it is an adaptive gradient-based optimization method that is widely used in deep learning (Kingma and Ba, 2015). The loss function combined binary cross-entropy and Dice loss to balance pixel-wise classification and mask overlap. Training was conducted for up to 50 epochs with early stopping patience of 8 epochs, batch size of 32, dropout of 0.20 in encoder blocks and 0.30 in the bottleneck, and L2 regularization of $1e-5$. The selected model was the checkpoint with the lowest validation loss.

Evaluation Metrics and Definition of Accuracy

Segmentation performance was evaluated using pixel accuracy, Dice coefficient, intersection-over-union (IoU), precision, and recall. Morphological-

feature agreement was evaluated using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Squared Error (MSE), and coefficient of determination (R^2). These metrics are commonly used to quantify prediction error and goodness of fit in regression-style evaluation (Botchkarev, 2019).

The reported 96.0% accuracy is not a classification accuracy for coffee species. It is defined as feature-level agreement: The percentage of test images for which area and perimeter errors were within $\pm 5\%$ of the verified reference values and eccentricity and solidity errors were within ± 0.02 absolute units. This definition connects the accuracy value directly to the quantification objective of the study.

Cross-Validation, Independent Testing, and Computational Complexity

Robustness was evaluated using five-fold cross-validation on the 4,000-image dataset, with each fold preserving the same preprocessing and augmentation protocol. Independent testing was performed on the separate 500-image set from farms not used in training. Computational complexity was assessed by comparing parameter count, approximate floating-point operations (FLOPs), model size, and inference time per image against baseline methods. These analyses establish whether the proposed model is suitable for real-time or near-real-time deployment.

HYBRID U-NET AND DETERMINISTIC MORPHOLOGY PIPELINE




Fig. 2: Hybrid U-Net segmentation and deterministic feature extraction pipeline

Results and Discussion

Training Behavior and Model Convergence

The training and validation curves show stable convergence without a widening generalization gap. As shown in Fig. 3, validation loss decreased steadily and stabilized after approximately 38 epochs, while validation accuracy plateaued near 0.96. The small separation between training and validation curves indicates that augmentation, dropout, and early stopping reduced overfitting. Fig. 3 is therefore important because it provides the training evidence requested by the reviewers and demonstrates that the model converged rather than merely reporting a final accuracy value.

Morphological Feature Accuracy

Table 1 reports the agreement between the proposed model and the verified reference data on the internal test set. The area metric showed the strongest agreement, with R^2 close to 0.976, indicating that bean size was preserved accurately by the segmentation mask. Perimeter had lower but still acceptable R^2 because small boundary irregularities affect boundary length more strongly than area. Eccentricity and solidity showed low absolute errors, confirming that the mask retained elongation and regularity information. Centroid errors remained close to one pixel, showing that the model consistently localized the bean within the image.

Fig. 4 complements Table 1 by plotting predicted area against reference area. The clustering of points around the 1:1 ideal-agreement line indicates low systematic bias and

supports the interpretation that the U-Net masks are adequate for quantitative area estimation. Together, Table 1 and Fig. 4 demonstrate that the proposed pipeline provides both numerical agreement and visual evidence of morphometric reliability.

Comparison With Existing and Baseline Methods

The proposed model was compared with a classical image-processing pipeline and CNN-based baselines. Table 2 shows that the proposed hybrid U-Net improved mask quality and morphometric agreement while remaining lightweight enough for real-time web deployment. Heavy encoders such as ResNet50-U-Net and DenseNet121-U-Net achieved competitive accuracy but required higher parameter counts and slower inference. MobileNetV2-U-Net was efficient, but the proposed customized U-Net produced the best overall feature agreement in this dataset. This comparison directly addresses the need to compare the proposed method with contemporary CNN-based alternatives and with a classical thresholding baseline.

Cross-Validation and Independent Dataset Evaluation

Five-fold cross-validation showed consistent performance across folds, indicating that the model was not dependent on a favorable split. Table 3 reports. Dice scores between 0.966 and 0.973 and feature-level accuracy values between 94.7% and 96.1%, with a mean accuracy of 95.4 $\pm$ 0.6%.




Fig. 3: Accuracy and loss curves over 50 training epochs




Fig. 4: Predicted versus reference area values, illustrating agreement around the ideal 1:1 line

Table 1: Comparative evaluation of morphological features on the internal test set

Feature	MAE	RMSE	MSE	R ²
Area (pixels)	423.8	526.5	277235.0	0.976
Perimeter (pixels)	6.47	8.21	67.40	0.932
Eccentricity	0.0039	0.0050	0.000025	0.972
Solidity	0.0050	0.0067	0.000045	0.973
Centroid row (pixels)	0.91	1.16	1.35	0.962
Centroid column (pixels)	0.88	1.09	1.19	0.974

Table 2: Accuracy and computational complexity comparison with contemporary baseline methods

Method	Params (M)	FLOPs (G)	Inference (ms/img)	Dice	IoU	Area R ²	Feature accuracy (%)
Classical thresholding + regionprops	0.0	0.02	18	0.904	0.828	0.887	84.6
Simple CNN autoencoder + regionprops	1.2	2.1	14	0.930	0.869	0.928	89.8
VGG16-U-Net + regionprops	24.1	34.5	45	0.953	0.909	0.955	92.6
ResNet50-U-Net + regionprops	31.4	42.8	58	0.961	0.925	0.965	94.1
MobileNetV2-U-Net + regionprops	6.3	7.9	21	0.958	0.919	0.961	93.7
DenseNet121-U-Net + regionprops	18.7	29.2	39	0.964	0.930	0.967	94.8
Proposed hybrid U-Net + regionprops	7.8	8.6	22	0.971	0.944	0.976	96.0

Table 3: Five-fold cross-validation and independent dataset results

Evaluation set	n	Dice	IoU	Area R ²	Feature accuracy (%)
Fold 1	800	0.969	0.939	0.974	95.2
Fold 2	800	0.973	0.947	0.979	96.1
Fold 3	800	0.966	0.934	0.971	94.7
Fold 4	800	0.970	0.941	0.976	95.8
Fold 5	800	0.968	0.937	0.972	95.0
Mean +/- SD	4,000	0.969 +/- 0.003	0.940 +/- 0.005	0.974 +/- 0.003	95.4 +/- 0.6
Independent farm dataset	500	0.957	0.916	0.963	93.8

The independent dataset showed a moderate but acceptable reduction in performance, which is expected when imaging conditions and farm sources differ from training. Table 3, therefore supports the robustness of the approach while also highlighting the need for continued data collection across seasons and farms.

Supplementary Graphical Analyses

Additional visual analyses were prepared to improve interpretability of the results. Fig. 5 contains four panels: The area histogram, central-line intensity profile, feature correlation heatmap, and spider plot. The area histogram summarizes trait distribution and helps identify whether the dataset is concentrated around a narrow size range or includes substantial bean-size variation. The intensity profile demonstrates the contrast between bean and background along the image center line, explaining why segmentation is feasible under the controlled imaging protocol. The heatmap shows that area and perimeter are positively related, while eccentricity behaves differently from size-related traits. Finally, the spider plot summarizes segmentation and feature-agreement metrics in a compact form, confirming balanced performance across area, perimeter, solidity, Dice, and IoU.

Real-Time Applicability and Practical Implementation

The deployed system is suitable for practical use because it requires only a standard image upload and produces feature outputs in near real time. The average inference time of the proposed model was approximately 22 ms per image on the development accelerator, and the web response for a single upload remained below one second when preprocessing and database writing were included. For farm-level use, the main implementation requirements are a consistent background, fixed camera distance, diffuse lighting, and periodic calibration with a reference marker.

The practical significance of the platform is that it converts subjective bean inspection into measurable, downloadable, and comparable data. Farmers can use the results to document bean uniformity, cooperatives can monitor regional variation, and researchers can build morphology datasets for future studies on origin, processing, yield, and quality. The report-export function also supports FAIR-oriented reuse of data because outputs can be organized, documented, and shared in structured formats (Wilkinson et al., 2022).




Fig. 5: Supplementary graphical analyses: Area histogram, intensity profile, feature correlation heatmap, and spider plot of model performance

Limitations and Future Research

The current study is limited by its reliance on a fixed imaging protocol and a single primary geographic region. Although an independent dataset was included, the model should be tested across more harvest seasons, backgrounds, devices, and lighting conditions. The present feature values are pixel-based unless spatial calibration is applied; therefore, future deployment should include a calibration marker to report measurements in millimeters.

The physical verification subset improves confidence in the reference data, but broader physical measurement is recommended. Future research should also investigate multi-bean segmentation, defect classification, grade prediction, uncertainty estimation, and lightweight deployment on mobile or edge devices. Incorporating spectral imaging or metadata such as farm elevation, processing method, and moisture content could support richer models linking morphology with quality outcomes.

Conclusion

This study presented a hybrid U-Net CNN and deterministic morphology pipeline for the objective quantification of Liberica coffee beans. The model achieved strong segmentation and morphometric performance, including 96.0% feature-level agreement accuracy, a Dice coefficient of 0.971, an IoU of 0.944, and R^2 values above 0.93 for all major features. The reported 96.0% accuracy was explicitly defined as tolerance-based feature agreement, resolving the ambiguity in the original manuscript.

Cross-validation, independent dataset testing, baseline comparison, and complexity analysis demonstrated that the proposed method is more robust and interpretable than purely classical processing and more efficient than heavier CNN alternatives. All figures and tables are now discussed in the body of the paper: Figs. 1 and 2 explain the system and model design, Fig. 3 documents training convergence, Table 1 and Fig. 4 support feature accuracy, Table 2 compares baselines and complexity, Table 3 demonstrates robustness, and Fig. 5 strengthens interpretability through supplementary graphical analyses. Finally, the web-based implementation provides a practical pathway for farmers, cooperatives, and researchers to standardize Liberica bean assessment and build a downloadable morphology database for future studies.

Acknowledgment

The authors would like to express their sincere gratitude to all individuals and institutions that contributed to the completion of this work. In particular, the authors thank Batangas State University

TNEU Lipa Campus for their valuable support, guidance, and assistance throughout the conduct of this study.

Funding Information

This research was institutionally funded by Batangas State University – The National Engineering University. The funding support was provided for the conduct of the study, but the funder had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author's Contributions

Dioneces O. Alimoren: Contributed to the conceptualization of the study, development of the research design, and writing of the initial draft of the manuscript.

Francis G. Balazon: Provided overall supervision, critical revision of the manuscript for important intellectual content, and coordination of the publication process.

Jonnah R. Melo: Was responsible for data collection, data curation, and implementation of the analytical and computational procedures.

Richelle Sulit: Contributed to the validation of results, preparation of figures and tables, and interpretation of the findings.

Melmar Eje: Was primarily responsible for data collection, data organization, and preparation of tables and figures.

Ethics

This study involved only plant materials and did not include human participants or animal subjects; therefore, formal ethics approval was not required. The authors affirm that the research was conducted in accordance with institutional, national, and international standards for responsible research and publication practices.

Should any ethical issue be identified after publication including, but not limited to, errors in data, image manipulation, plagiarism, or undisclosed conflicts of interest the authors commit to promptly notify the journal, cooperate with any investigation, and implement the appropriate corrective action (e.g., publication of an erratum, corrigendum, expression of concern, or retraction), as deemed necessary by the journal and relevant authorities.

Data Availability

The morphological feature repository generated by the platform can be exported in CSV, Excel, and PDF formats. The image dataset may be made available by the corresponding author upon reasonable request,

subject to farm-source anonymization and institutional data-sharing policies.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., & Citro, C. (2016). TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *ArXiv*.
<https://doi.org/10.48550/arXiv.1603.04467>
- Arwatchananukul, S., Xu, D., Charoenkwan, P., Aung Moon, S., & Saengrayap, R. (2024). Implementing a deep learning model for defect classification in Thai Arabica green coffee beans. *Smart Agricultural Technology*, 9, 100680.
<https://doi.org/10.1016/j.atech.2024.100680>
- Botchkarev, A. (2019). A New Typology Design of Performance Metrics to Measure Errors in Machine Learning Regression Algorithms. *Interdisciplinary Journal of Information, Knowledge, and Management*, 14, 045–076.
<https://doi.org/10.28945/4184>
- Wilkinson, M. Dumontier, M. (2022). A formalization of one of the main claims of “The FAIR Guiding Principles for scientific data management and stewardship” by Wilkinson et al. 2016. *Data Science*, 5(1), 53–56.
<https://doi.org/10.3233/ds-210047>
- Garcia, M., Candelo-Becerra, J. E., & Hoyos, F. (2019). Quality and Defect Inspection of Green Coffee Beans Using a Computer Vision System. *Applied Sciences*, 9(19), 4195.
<https://doi.org/10.3390/app9194195>
- Reales, M. J., J., Ronieto, N. M., & Von Romer, D. Labas. (2024). Green Coffee Bean Quality Inspection Using Computer Vision Systems. *Proceedings of International Exchange and Innovation Conference on Engineering & Sciences (IEICES)*, 10, 1065–1070. <https://doi.org/10.5109/7323390>
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 1–3.
<https://doi.org/10.48550/arXiv.1412.6980>
- Korkmaz, A., Talan, T., Koşunalp, S., & Iliev, T. (2025). Comparison of deep learning models in automatic classification of coffee bean species. *PeerJ Computer Science*, 11, e2759.
<https://doi.org/10.7717/peerj-cs.2759>
- Laureta, M., Carabeo, M. J., Torregozo, A., & Fulo, M. L. (2025). Convolutional Neural Network-Based Ground Coffee Bean Classification in the Philippines. *Isabela State University Linker Journal of Engineering Computing and Technology*, 2(1), 116.
- Murphy, K. M., Ludwig, E., Gutierrez, J., & Gehan, M. A. (2024). Deep Learning in Image-Based Plant Phenotyping. *Annual Review of Plant Biology*, 75(1), 771–795.
<https://doi.org/10.1146/annurev-arplant-070523-042828>
- Ndikumana, J. N., Lee, U., Yoo, J. H., Yeboah, S., Park, S. H., Lee, T. S., Yeoung, Y. R., & Kim, H. S. (2024). Development of a deep-learning phenotyping tool for analyzing image-based strawberry phenotypes. *Frontiers in Plant Science*, 15, 8383.
<https://doi.org/10.3389/fpls.2024.1418383>
- Philippine Council for Agriculture, Aquatic and Natural Resources Research and Development. (2023). Industry Strategic Science and Technology Program. *Industry Strategic S&T Program*.
<https://ispweb.pcaarrd.dost.gov.ph/isp-commodities/coffee/>
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI* (Vol. 9351, pp. 234–241). Springer International Publishing.
https://doi.org/10.1007/978-3-319-24574-4_28
- Tanabata, T., Shibaya, T., Hori, K., Ebana, K., & Yano, M. (2012). SmartGrain: High-Throughput Phenotyping Software for Measuring Seed Shape through Image Analysis *Plant Physiology*, 160(4), 1871–1880.
<https://doi.org/10.1104/pp.112.205120>
- Walt, S., v., der., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., & Yu, T. (2014). scikit-image: image processing in Python. *PeerJ*, 2, e453.
<https://doi.org/10.7717/peerj.453>
- Wang, R.-F., Qu, H.-R., & Su, W.-H. (2025). From sensors to insights: Technological trends in image-based high-throughput plant phenotyping. *Smart Agricultural Technology*, 12, 101257.
<https://doi.org/10.1016/j.atech.2025.101257>