Estimation of leaf dimensions through image segmentation using a new annotated dataset

Gabriela Singulani Marques[1], João Victor Corrêa de Almeida[1], Karla Gabriele Florentino da Silva[2], Artur Ferreira de Castro[1], João Pedro Lemos Guadalupe[1], Vinícius Macuco Leite Carvalhaes[1], Paulo Victor de Magalhães Rozatto[1], Luiz Maurílio da Silva Maciel[1], Marcelo Bernardes Vieira[1], Saulo Moraes Villela[1]

[1]Departamento de Ciência da Computação – Universidade Federal de Juiz de Fora
Juiz de Fora – MG – Brasil

[2]Instituto de Computação – Universidade Estadual de Campinas
Campinas – SP – Brasil

{gabriela.singulani,joaovictor.correa,artur.castro,guadalupe.joao,
macuco.vinicius,paulo.rozatto}@estudante.ufjf.br
k272454@dac.unicamp.br
{luiz.maciel,marcelo.bernardes,saulo.moraes}@ufjf.br

Abstract. *The leaf is the plant organ responsible for photosynthesis, and its morphology provides important insights. Traditional methods for estimating leaf morphology are labor-intensive, while dedicated measurement devices present high financial costs. Computer vision–based approaches emerge as a cost-effective solution. In this work, we introduce a dataset of leaf images, each accompanied by a reference square pattern. The dataset includes high-quality segmentation masks and ground-truth leaf dimensions. We applied segmentation networks on this dataset and used the resulting masks to estimate leaf dimensions using an image processing method. The experiments establish baseline results for segmentation and leaf-dimension estimation on the proposed dataset.*

1. Introduction

The leaf is the plant organ responsible for photosynthesis. Through this process, plants convert CO2 into energy, ensuring their nutrition [Roth-Nebelsick and Krause 2023]. Leaves exhibit a wide variety of shapes depending on the species. The study of leaf morphology is essential for understanding several processes, including biomass production [Fan et al. 2024], plant productivity and growth [Tong et al. 2023], and environmental pollution [Zhang et al. 2025].

Methods for estimating leaf morphological parameters (e.g., area, perimeter, width, and length) include high-cost scanners and manual approaches [Pandey and Singh 2011]. However, high-cost scanners are often inaccessible to many researchers, while manual methods are labor-intensive and time-consuming.

In recent years, methods based on computer vision and image processing have emerged. Some approaches rely solely on conventional image processing techniques [Tech et al. 2018], which are often sensitive to noise and variations in image quality. More recently, deep learning–based methods have been developed. In particular, semantic seg-

mentation networks have been widely applied to detect leaves and support the estimation of morphological parameters [Giang and Ryoo 2024, Richardson et al. 2023].

Annotated datasets are essential for deep learning–based approaches. In particular, models based on semantic segmentation networks typically require datasets with accurately annotated pixel-level masks. Furthermore, for leaf dimension estimation tasks using computer vision, it is necessary to have images for which the ground-truth dimensions are known. Despite the increasing use of deep learning for leaf analysis, there is a scarcity of datasets providing multi-species leaf images together with both pixel-level segmentation masks and reference measurements of leaf dimensions.

In this work, we introduce a comprehensive dataset comprising high-quality segmentation mask annotations and ground-truth leaf dimensions. The dataset includes images from multiple species, each accompanied by a square pattern used for scale estimation. We provide precise segmentation masks as well as leaf dimensions obtained through careful manual estimates. We evaluate the application of semantic segmentation networks on the images in our dataset. The resulting masks are then used to estimate leaf dimensions through a straightforward approach.

Therefore, the main contribution of this work is a novel dataset comprising 1,088 leaf images from 13 species, annotated with high-quality segmentation masks and ground-truth leaf dimensions. To demonstrate its use and establish reproducible baselines, we evaluate four segmentation models and a mask-based image-processing pipeline for leaf-dimension estimation.

2. Related Work

Datasets of leaf images with annotated segmentation masks remain relatively scarce in the literature. An example is the work of [Silva et al. 2025b], which introduced a dataset comprising 6,981 images of bean leaves captured under uncontrolled environmental conditions. Similar to the dataset proposed in the present work, the authors also provide estimates of leaf dimensions. However, their dataset is limited to a single species, whereas our dataset encompasses 13 different species. Another relevant example is the dataset proposed by [Segreto et al. 2026], which contains 640 images with 7,221 soybean and 5,190 cotton leaves annotated at the instance level with bounding boxes and segmentation masks. Despite the large number of annotated leaf instances, the dataset does not provide ground-truth leaf dimensions, limiting its applicability for the development and evaluation of leaf dimension estimation methods.

In recent years, studies have proposed methods for estimating leaf dimensions from digital images using computer vision techniques. [Tech et al. 2018] proposed a destructive approach based on conventional image processing techniques. Their method applies image thresholding followed by leaf area estimation through pixel counting, using a square reference marker positioned at a fixed location on the acquisition table. A similar approach was proposed by [Silva et al. 2023], in which leaf dimensions are estimated using a simple image processing pipeline. The method performs image thresholding, filters noisy contours, and identifies the square calibration pattern based on its geometric properties. However, both approaches are sensitive to image noise and illumination artifacts, such as shadows, which may compromise segmentation quality and, consequently, dimension estimation accuracy.

Some studies have employed neural networks to support leaf dimension estimation. [Richardson et al. 2023] introduced the Phenobot system, which estimates leaf area through image segmentation using a red square calibration pattern as a scale reference. However, the system relies on a controlled acquisition setup with a fixed structure. [Giang and Ryoo 2024] combined RGB-D images with semantic segmentation networks to estimate leaf area. Their approach is based on counting points on the leaf surface and depends on a controlled acquisition setup involving two fixed cameras. Another study by [Silva et al. 2025a] estimated the area of bean leaves directly on the plant using an adapted semantic segmentation network. In this approach, the network directly predicts an area map from the input image. While these previous studies focus exclusively on leaf area estimation, the method proposed in the present work enables the estimation of multiple morphological parameters, including area, perimeter, width, and length.

3. Dataset

This section describes the proposed dataset for leaf segmentation and dimension estimation. The dataset is publicly available at the following link: https://data.mendeley.com/datasets/s654c5wvkh/1

3.1. Leaf Collection and Image Acquisition

Our dataset comprises images from 13 species: bean, bell pepper, castor bean, coffee, corn, cotton, eucalyptus, guava, lemon, orange, pumpkin, soybean, and tomato. Leaves were detached from the plant in the morning period (before 9 a.m.) to ensure adequate turgidity. We collected healthy leaves, without visible signs of nutritional deficiency, pathogen infection, or disease. Image acquisition and manual dimension estimation were performed on the same day to preserve leaf integrity and prevent degradation.

For image acquisition, leaves were placed on a table covered with white paper to ensure a uniform and controlled background. A black square of known dimensions was positioned alongside the leaves. The size of this pattern is provided in the dataset to serve as a scale reference. A glass plate was placed over the samples to keep the leaves as flat as possible. Images were captured from a perpendicular viewpoint at varying distances. Multiple leaves were included in some images to reflect the practical scenario of analyzing batches of leaves simultaneously.

Data acquisition was performed using five smartphones, a digital camera, and a scanner, ensuring variability in imaging conditions. The dataset includes images containing both single and multiple leaves. Figure 1a illustrates an example of a captured image, while Table 1 summarizes the acquisition settings.



(a) (b)

Figure 1. Dataset image sample. (a) Original image. (b) Segmentation mask.

Table 1. Capture configurations of the proposed dataset.

Device/resolution(w,h)	OneVision/(4000, 3000), MotoZ2/(4032, 3024), MotoC/(1920, 2560), LGK430TV/(4160, 3120), Scanner/(1240, 1755), Digital Camera/(2976, 1984), LGK10/(4160, 3120)
Pattern Side (cm)	1, 2, 3, 4, 5, 6
Capture Distance (cm)	12, 15, 16, 20, 30, 40, 50, 70

3.2. Manual Estimation of Leaf Dimensions and Image Annotation

We employed a manual paper-weight–based method [Pandey and Singh 2011] to estimate leaf dimensions. Initially, the leaf contour was traced onto a sheet of paper, which was subsequently cut and weighed using a precision balance. Given the known paper grammage, the leaf area was then estimated.

For the perimeter estimate, a flexible string was carefully aligned along the contour of the paper cutout and subsequently measured with a ruler. Leaf length and width were directly estimated on the paper contour using a ruler.

For each image in the dataset, segmentation masks were annotated considering three classes (background, leaf, and square). Figure 1b presents an example of an annotated mask. The annotation process involved a two-stage review procedure. Initially, each image was semi-automatically annotated by a member of the research team using an annotation software developed by [Rozatto et al. 2025]. Subsequently, a second team member carefully reviewed the annotations and performed the necessary refinements.

We provide an XML (Extensible Markup Language) file containing the annotations for each image. The annotations include metadata related to the acquisition conditions (device, scale pattern size, capture distance, and image resolution), as well as leaf-specific information (species and index) and the discretized mask representations (normalized by the image height). Table 2 summarizes the general structure of the XML file.

Table 2. Description of the annotation tags of the dataset.

Tag	Description	Type
<image-name>	Name of the image.	Text
<capture-distance>	Distance of the image capture in centimeters.	Integer
<pattern-side>	Length of the side of the pattern of the image in centimeters.	Integer
<device>	Name of the device used to capture the image.	Text
<leaf-quantity>	Quantity of leaves in the image.	Integer
<image-size>(subtags <height>, <width>, <depth>)	Image dimensions.	Integer

For each leaf instance, the XML file includes a <leaf> tag, whose subtags are detailed in Table 3. In some cases, certain leaf dimensions could not be reliably estimated.

Instead of discarding these samples, we retained them in the dataset and indicated the missing values in the XML file using the label "NAva". For dimensions that were not available, we did not compute errors in the dataset evaluation. For the scale pattern, the XML file includes a <pattern> tag, with subtags <bbox> and <polygon>, defined analogously to those of the <leaf> tag.

Table 3. Description of the subtags of <leaf> tag.

Tag	Description	Type
<index>	Index of the leaf (based on its species).	Integer
<specie>	Species of the leaf.	Text
<dimensions>(subtags <area>,<perimeter>, <length>, <width>)	Known dimensions of the leaf.	Real
<bbox>	Bounding box of the leaf, composed by the tags <x>and <y> representing the upper-left corner and the tags <width> and <height> indicating the bounding box dimensions. All values are normalized by the image height.	Real
<polygon>	A list of points as tags <x1>, <y1>, <x2>, <y2>, ..., <xn>, <yn>representing the position of the points of the leaf contour, normalized by the image height.	Real

3.3. Dataset Organization

Our dataset comprises 1,088 images, of which 1,074 contain leaves from a single species. Table 4 presents the distribution of images and leaf instances across species for the images containing a single species. A significant class imbalance is observed, ranging from 202 images containing 30 corn leaves to only 20 images comprising 2 castor bean leaves. The table also reports the total number of leaf instances per species across all images. The dataset includes 14 images containing leaves from multiple species within the same image.

To facilitate the use of the dataset, we propose a 5-fold cross-validation protocol. The folds were constructed according to the following criteria:

- Images containing a single leaf were assigned such that all images of the same leaf were placed within the same fold;
- Images containing multiple leaves were distributed to ensure that samples of the same species were grouped within the same fold.

This strategy is designed to prevent data leakage in images containing single leaves and to mitigate potential leakage in images with multiple leaves. Figure 2 presents the distribution of images across the folds. Due to the criteria adopted for fold construction, the folds do not contain an equal number of images.

Table 4. Image distribution by species in the dataset.

Species	Leaves	Images	Instances
Bean	18	63	338
Bell Pepper	2	24	32
Castor Bean	2	20	24
Coffee	27	146	304
Corn	30	202	412
Cotton	11	61	244
Eucalyptus	31	162	318
Guava	3	29	29
Lemon	2	24	24
Orange	20	80	232
Pumpkin	2	23	31
Soybean	30	170	472
Tomato	7	70	96
Total	**185**	**1074**	**2556**

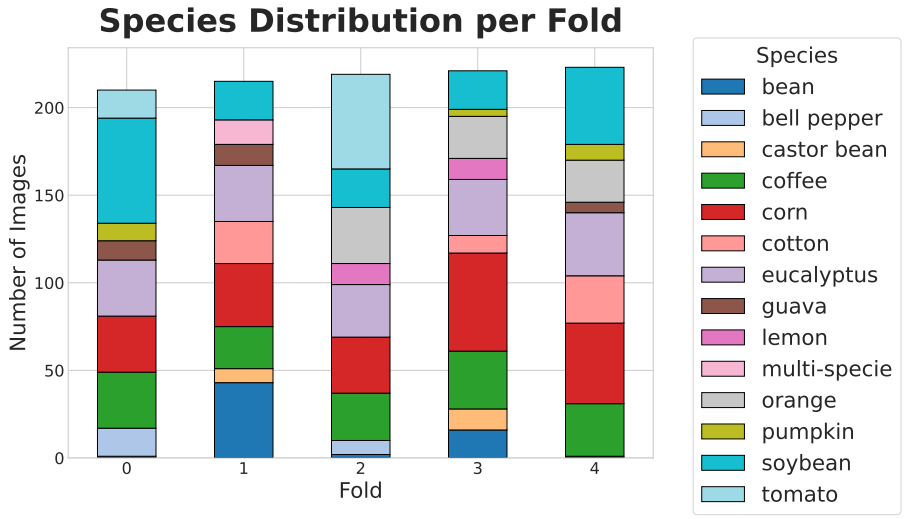


Figure 2. Images distribution per species in the proposed 5-fold protocol.

4. Semantic Segmentation Networks

We applied four models to segment the images of our dataset: YOLOv8, YOLOv11, MobileSeg, and DeepLabv3+. The You Only Look Once (YOLO) [Redmon et al. 2016] family of models performs object detection in a single forward pass, making it adequate for real-time tasks. In this work, two versions were evaluated: YOLOv8 and YOLOv11, both of which support semantic segmentation. The YOLOv8 version built upon previous versions by using a modern backbone and neck, which enhances feature extraction capabilities. Additionally, it adopts an anchor-free split head, thereby improving computational efficiency. YOLOv11 further advances the YOLO framework by incorporating an optimized backbone and neck architecture, designed to improve both feature representation and overall task performance.

PP-MobileSeg [Tang et al. 2023] is an ultra-lightweight semantic network optimized for mobile devices. To handle different scales and refine edges, the architecture employs the Aggregate Attention Module (AAM), which filters out irrelevant backgrounds, and the Valid Interpolate Module (VIM), which focuses on the present classes for fast and computationally low-cost upsampling. DeepLabv3+ [Chen et al. 2018] adopts an encoder-decoder structure, where the encoder employs the Atrous Spatial Pyramid Pooling (ASPP) module to capture multi-scale context through atrous convolutions.

5. Estimation of leaf dimensions

Although annotations for both square and leaf objects are available, treating all objects as a single class yielded superior segmentation performance. Thus, the semantic segmentation models receive an input image I and produce a mask M, in which each pixel is classified either as an object of interest (leaf or square) or as background. To estimate the leaf dimensions, we need to identify the square scale pattern based on its geometry. Initially, we estimated the area of each contour in the image using the Green formula. We discarded contours larger than 10^{10} pixels and smaller than 300 pixels, as they typically correspond to noise. For each remaining contour C_k in the mask M, we computed the best-fitting bounding rectangle. Subsequently, we selected the contours whose rectangle exhibited an aspect ratio smaller than 1.5 as candidate contours. Among these candidates, the square pattern was identified as the contour whose bounding rectangle has the highest filling ratio. This criterion was based on the assumption that the square pattern exhibits an approximately rectangular shape in the image.

After identifying the square pattern, we have contours of each leaf L_i and the square S. Let a_i denote the area of the leaf L_i in pixels, a_s the area of the square S in pixels, and a_s^r the known area of the square scale pattern. The area a_i^r of the leaf can be estimated as follows:

$$a_i^r = a_s^r \frac{a_i}{a_s}. \tag{1}$$

For each leaf, we computed the perimeter of the contour in pixels, denoted by p_i, along with the perimeter of the contour of the square, denoted by p_s. Given the known real perimeter of the square, p_s^r, we estimated the leaf perimeter, p_i^r, as:

$$p_i^r = p_s^r \frac{p_i}{p_s}. \tag{2}$$

For length and width estimation, the first step of our method consisted of aligning the major axis of the leaf with the horizontal direction. To achieve this, we applied the Principal Component Analysis (PCA) to determine the principal orientation of each contour and rotated it according to this direction.

After contour alignment, we estimated the leaf length and width using the contour bounding box. We defined the length in pixels, denoted by l_i, as the width of the bounding box and the width, denoted by w_i, as the height of the bounding box. For the square pattern, we calculated the side length in pixels as one-fourth of the perimeter. Given the known real side length of the square pattern, s_s^r, we estimated the leaf length, l_i^r, and width, w_i^r, as:

$$l_i^r = s_s^r \frac{l_i}{l_s}, \tag{3}$$

$$w_i^r = s_s^r \frac{w_i}{w_s}. \tag{4}$$

It is important to note that our method assumes that the largest dimension of the leaf corresponds to its length. Although this assumption holds for many leaf shapes, it may lead to inaccurate estimations in certain cases. This constitutes a limitation of our method and should be addressed in future work.

6. Experiments and Results

In this section, we describe the main experiments and results achieved. We include the results for the segmentation problem and for the leaf dimension estimation.

6.1. Segmentation Results

We evaluated the performance of the four semantic segmentation models on our proposed dataset. Table 5 describes the main training configuration for each model. In the table, "LR" denotes learning rate, "*" indicates that YOLOv11 employs a native backbone, and "**" indicates that PP-MobileSeg receives the image in the original resolution and performs random crops of 768 $\times$ 768.

Table 5. Training configuration for each evaluated network.

Model	LR	Scheduler	Batch	Epochs	Backbone	Input
YOLOv8n	0.01	Linear	16	up to 200	CSPDarknet53	640 $\times$ 640
YOLOv11n	0.001	Linear	16	up to 200	*	640 $\times$ 640
PP-MobileSeg	0.0004	Poly	9	up to 100	StrideFormer	**
DeepLabv3+	0.0003	Poly	2	70	Xception	512$\times$ 512

Although we have the annotation of the square and leaf objects, considering all objects as a single class achieved better results. We evaluated the performance using the mean Intersection over Union (mIoU) metric, commonly used to evaluate semantic segmentation models. We ran the complete cross-validation, analyzed the performance for each test set, and combined them to compute the mIoU on the complete dataset. For each training, we adopt a fold for validation following the pattern: when fold 1 is test, then fold 2 is validation; when fold 2 is test, then fold 3 is validation; and so on. Table 6 presents the results for each model.

Table 6. Results for mIoU for each model evaluated.

Model	mIoU	Model	mIoU
YOLOv8	0.9089	PP-MobileSeg	0.9863
YOLOv11	0.9107	DeepLabv3+	0.9907

We observed good performance across all networks on the leaf segmentation task. In particular, DeepLabv3+ achieved an mIoU greater than 0.99, segmenting objects with high quality. However, some limitations were observed. There are problems regarding noise present in the images. For example, in some images, the table where the capture was made became apparent on the borders. This is not expected in an ideal scenario for our method. On the other hand, in general, the networks were able to handle well variations in lighting and shadow.

6.2. Leaf Dimension Estimation Results

For each leaf on each image, we estimated the dimensions using the method described in Section 5. As in the segmentation experiments, we estimated the dimensions using only the test sets of the cross-validation protocol. Consequently, we obtained the dimensions of all the images in the dataset. To evaluate the results, we compared the dimensions with the corresponding known dimensions manually annotated. We evaluated the results based on the relative error rate (RER), defined as $RER(\%) = \frac{|v_{est} - v_{gt}|}{v_{gt}} \cdot 100$, where v_{est} denotes the estimated value and v_{gt} represents the corresponding ground truth value. In our experiments, we assumed the dimensions estimated by using the manual procedure as the ground truth.

Table 7 presents the mean and standard deviation of the RER for the dimensions estimated using the segmentation outputs of each network. The mean RER rates are between 10% and 15% for most dimensions and evaluated networks. In particular, PP-MobileSeg achieved the smallest mean RER for most dimensions, whereas YOLOv8 obtained the lowest mean RER for area estimation. However, we observed high standard deviation values for most results. These high deviations are primarily caused by the presence of outliers, which also contribute to an increase in the mean RER values. Such outliers mainly arise from incorrect detection of the scale pattern in some images. Another reason is that some leaves are positioned very close together in certain images. As a consequence, adjacent leaves may be merged during segmentation, resulting in erroneous estimates of dimensions. Nevertheless, only a small subset of images is affected by these issues.

Table 7. RER (mean $\pm$ standard deviation) for dimension and model.

Model	Area	Length	Perimeter	Width
YOLOv8	**13.74 $\pm$ 14.73**	11.75 $\pm$ 12.24	12.45 $\pm$ 13.88	32.80 $\pm$ 82.68
YOLOv11	14.23 $\pm$ 15.14	12.01 $\pm$ 12.72	12.87 $\pm$ 14.86	35.20 $\pm$ 88.19
PP-MobileSeg	17.43 $\pm$ 136.83	**10.70 $\pm$ 25.92**	**12.40 $\pm$ 37.14**	**12.20 $\pm$ 53.88**
DeepLabv3+	18.91 $\pm$ 86.71	11.83 $\pm$ 21.22	21.96 $\pm$ 42.06	13.54 $\pm$ 31.18

Table 8 presents the distribution of RER values below different thresholds to provide a more detailed analysis of the results and reduce the influence of outliers. The results further confirm the superior performance of PP-MobileSeg. Across all evaluated dimensions, at least 37.0% of the estimations achieved an error below 5%, while at least 62.8% presented errors smaller than 10%. Furthermore, for this network, only approximately 10% of the estimations exhibited errors greater than 20%. DeepLabv3+ also achieved competitive performance, although its results remained below those obtained by PP-MobileSeg. The YOLO-based models exhibited difficulties in the leaf dimension estimation task. These limitations can be attributed to irregularities in the generated segmentation masks, which directly affect the subsequent estimation process. Since the YOLO architecture was originally designed for object detection rather than semantic segmentation, it may present certain limitations when applied to pixel-level segmentation tasks.

Figure 3 presents scatter plots of the leaf dimensions for each evaluated network. In addition, each plot includes a fitted trend line representing the overall estimation tendency. We set the maximum value in the y-axis as the maximum expected (ground-truth)

Table 8. RER distribution for each model evaluated and specific thresholds.

	Area				Length			
Model	≤ 5	≤ 10	≤ 15	≤ 20	≤ 5	≤ 10	≤ 15	≤ 20
YOLOv8	28.7	49.4	67.2	79.3	31.7	59.3	74.9	84.1
YOLOv11	27.0	48.9	65.5	77.4	30.9	57.5	74.1	83.8
PP-MobileSeg	**37.7**	**62.8**	**79.3**	**87.0**	**40.1**	**71.8**	**88.2**	**92.9**
DeepLabv3+	34.2	60.0	75.8	82.5	32.4	60.2	75.2	82.4
	Perimeter				**Width**			
YOLOv8	29.6	54.0	72.6	83.0	30.2	54.7	69.7	80.0
YOLOv11	29.1	52.7	72.3	82.9	29.1	53.8	68.5	78.0
PP-MobileSeg	**37.0**	**68.3**	**86.8**	**92.4**	**48.1**	**70.9**	**81.3**	**89.4**
DeepLabv3+	17.7	34.3	49.2	64.3	43.7	67.8	77.8	84.8

value on the x-axis. Consequently, some outliers are not displayed in the graphics. We note that most points lie close to the identity line, indicating good agreement between estimated and ground-truth values. However, some values are overestimated or underestimated. Overall, the scatter plots for PP-MobileSeg and DeepLabv3+ show better agreement between predicted and reference values compared to the YOLO-based models, which is consistent with the RER results.

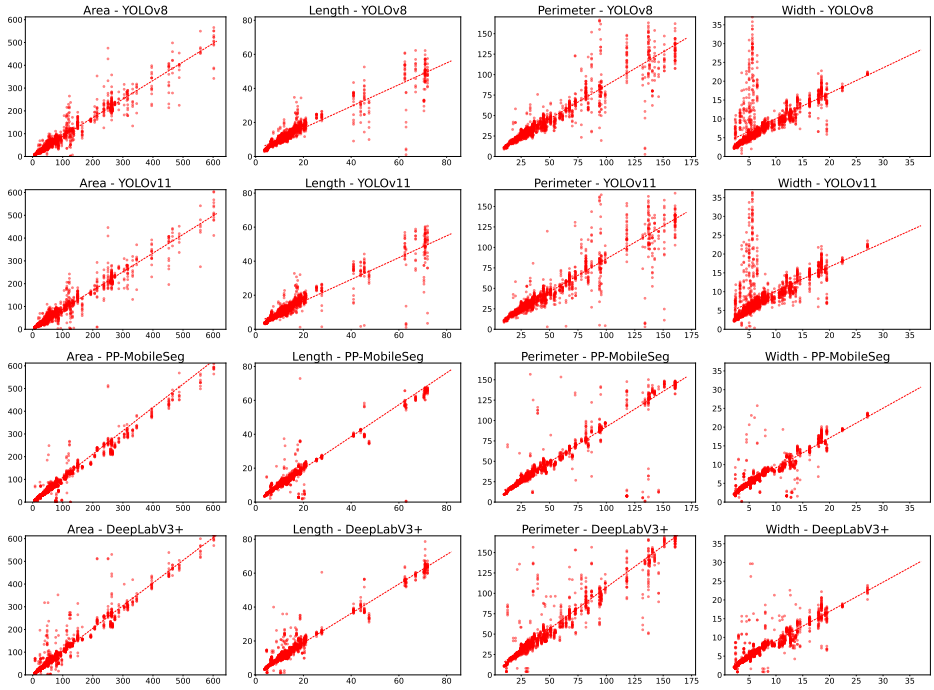


Figure 3. Scatter plots for each dimension and evaluated network. The x-axis represents the ground-truth values and the y-axis the corresponding estimated values.

We observed that some outliers produced by the proposed method were due to leaves that were fused in the segmentation masks. These leaves had been positioned very close to one another during image acquisition. To investigate the impact of this issue,

we identified all leaves fused in the segmentation output of at least one network, and we conducted an additional analysis that excluded them. We excluded 159 leaves from 36 images. Table 9 presents the results of this analysis. An improvement can be observed for most networks and evaluated dimensions, confirming the influence of fused leaves on the estimation errors. This issue could be mitigated by increasing the distance between leaves during image acquisition or by employing instance segmentation models capable of separating individual leaves.

Table 9. RER (mean ± standard deviation) for dimension and model, excluding fused leaves.

Model	Area	Length	Perimeter	Width
YOLOv8	**12.62 ± 11.42**	10.92 ± 10.19	11.14 ± 9.94	31.27 ± 80.74
YOLOv11	13.51 ± 11.91	14.04 ± 10.80	11.31 ± 10.42	21.45 ± 18.30
PP-MobileSeg	13.20 ± 66.92	**9.42 ± 16.87**	**9.89 ± 18.05**	**9.81 ± 19.13**
DeepLabv3+	13.29 ± 73.44	9.66 ± 13.89	16.83 ± 16.91	9.89 ± 14.61

7. Conclusion

In this work, we introduced a dataset composed of multi-species leaf images, accompanied by a square scale pattern. Our dataset provides annotations including high-quality segmentation masks and leaf dimensions (area, perimeter, length, and width) obtained through a manual measurement methodology. We evaluated four semantic segmentation networks on the proposed dataset, and all achieved mIoU values greater than 0.90 in the segmentation task.

The segmentation outputs were subsequently used as input to a simple method for estimating leaf dimensions. The results indicate the potential of the proposed dataset for supporting the development and evaluation of leaf dimension estimation methods. However, some outliers were identified, indicating opportunities for further improvements in the estimation process.

As future work, we intend to evaluate additional semantic segmentation networks and training configurations. Furthermore, we can propose improvements in our dimension estimation methodology, such as better strategies for identifying the length and width directions.

References

Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H. (2018). Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, Munich, Germany. ECCV.

Fan, X. et al. (2024). Plant functional traits modulate the effects of soil acidification on above- and belowground biomass. *Biogeosciences*, 21:2641–2656.

Giang, T. T. H. and Ryoo, Y.-J. (2024). Sweet pepper leaf area estimation using semantic 3d point clouds based on semantic segmentation neural network. *AgriEngineering*, 6(1):645–656.

Pandey, S. and Singh, H. (2011). A simple, cost-effective method for leaf area estimation. *Journal of botany*, 2011(1):658240.

Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2016). You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788.

Richardson, G. A., Lohani, H. K., Potnuru, C., Donepudi, L. P., and Pankajakshan, P. (2023). Phenobot: an automated system for leaf area analysis using deep learning. *Planta*, 257(2):36.

Roth-Nebelsick, A. and Krause, M. (2023). The plant leaf: A biomimetic resource for multifunctional and economic design. *Biomimetics*, 8(2):145.

Rozatto, P. V. M., Maciel, L. M. S., Vieira, M. B., and Villela, S. M. (2025). Development of an artificial intelligence-aided software for annotating image datasets. In *Brazilian Symposium on Information Systems (SBSI)*, pages 105–114, Recife, Brazil.

Segreto, T. H., Negri, J. D., Polegato, P. H., Pinheiro, J. M. H., Godoy, R. V., and Becker, M. (2026). A leaf-level dataset for soybean-cotton detection and segmentation. *Scientific Data*.

Silva, K. G. F., Maciel, L. M. S., Villela, S. M., Pedrini, H., Morais, L. E., and Vieira, M. B. (2025a). Non-destructive leaf area estimation based on a semantic segmentation deep neural network. *Computers and Electronics in Agriculture*, 237:110551.

Silva, K. G. F., Moreira, J. M., Calixto, G. B., Maciel, L. M. S., Miranda, M. A., and Morais, L. E. (2023). A simple and low-cost method for leaf surface dimension estimation based on digital images. In *12th Brazilian Conference on Intelligent Systems, BRACIS 2023*, pages 146–161. Springer.

Silva, K. G. F., Rozatto, P. V. M., Sousa, K. O., Hudson, L. D., Meyer, A. W. S., Silva, A. F., Mendes, I. T., Borges, A. R., Morais, L. E., Maciel, L. M. S., Villela, S. M., Pedrini, H., and Vieira, M. B. (2025b). Bean leaf image dataset annotated with leaf dimensions, segmentation masks, and camera calibration. *Data in Brief*, 59:111328.

Tang, S., Sun, T., Peng, J., Chen, G., Hao, Y., Lin, M., Xiao, Z., You, J., and Liu, Y. (2023). Pp-mobileseg: Explore the fast and accurate semantic segmentation model on mobile devices. *arXiv preprint arXiv:2304.05152*.

Tech, A. R. B., Silva, A. L. C., Meira, L. A., Oliveira, M. E., and Pereira, L. E. T. (2018). Methods of image acquisition and software development for leaf area measurements in pastures. *Computers and Electronics in Agriculture*, 153:278–284.

Tong, R., Yang, X., Wang, Q., Li, L., Shi, Y., Mu, C., and Wang, J. (2023). The shift in key functional traits caused by precipitation under nitrogen and phosphorus deposition drives biomass change in leymus chinensis. *Plants*, 12(9):1781.

Zhang, Y. et al. (2025). Leaf ecological traits (morphology and gas exchange) and polycyclic aromatic hydrocarbons concentrations in shrubs and trees: A meta-analysis approach. *Environmental Pollution*, 364:125337.