LatexLens: An Interactive AI-Powered Annotation and Quantification Framework for Laticifers in Plant Microscopy

Sergio Madrid Pérez[1] , **Sofía Martínez Pastor**[1] , **Francisco J. Escaray**[2] , **Pablo Vera**[2] , **Cèsar Ferri**[1] , **Carlos Monserrat**[1]

[1]Valencian Research Institute for Artificial Intelligence, Universitat Politècnica de València, València, Spain

[2]Instituto de Biología Molecular y Celular de Plantas, Universidad Politécnica de Valencia-Consejo Superior de Investigaciones Científicas, València, Spain

smadper@upv.edu.es, smarpas@ade.upv.es, fjescmaz@ibmcp.upv.es, vera@ibmcp.upv.es, cferri@dsic.upv.es, cmonserr@dsic.upv.es

Abstract

The analysis of laticifers in plant tissues is crucial for understanding latex production and plant defence mechanisms. However, the segmentation of these microscopic structures is hindered by their intricate topology, low contrast, and high-resolution imaging requirements. While Deep Learning models have shown promise, the gap between raw model inference and usable biological data remains significant due to the lack of accessible interfaces for domain experts. In this paper, we present LatexLens, a comprehensive desktop application built upon the napari ecosystem. The tool integrates a U-Net-based inference engine, domain-specific preprocessing (CLAHE), and a Human-in-the-Loop (HITL) workflow for mask refinement. Furthermore, we implement automated biological quantification directly within the GUI, focusing on metrics botanists use in practice: laticifer area fraction within the tissue and transect-based intersection counts that estimate network density. This focus bridges the gap between computer vision outputs and reproducible botanical measurements.

Demo Video: https://youtu.be/iH57-hDS740

1 Introduction

Laticifers are specialised internal secretory structures found in various plant species, responsible for the production and storage of latex [Fahn, 1979]. Understanding the morphology and distribution of these cells is essential for applications ranging from plant physiology to the industrial extraction of natural rubber and pharmaceutical compounds [Agrawal and Konno, 2009; Hagel *et al.*, 2008]. However, analysing these structures in microscopy images poses significant computer vision challenges. Laticifers often appear as thin, elongated, and anastomosing networks embedded within complex tissue backgrounds, often suffering from low contrast and noise.

Recent advancements in Deep Learning, particularly Convolutional Neural Networks (CNNs) [LeCun *et al.*, 1989], such as U-Net [Ronneberger *et al.*, 2015], have revolutionised biomedical image segmentation. However, a persistent bottleneck in deploying these models is the lack of user-friendly interfaces that enable biologists to leverage these algorithms effectively. Furthermore, a segmentation mask itself is not the end goal; biological insights depend on derived calculations such as the laticifer area fraction within the image, and line-intersection counts that reflect cell density.

To address these challenges, we propose LatexLens, an interactive annotation and quantification tool. Built on the napari [Chiu *et al.*, 2022] multi-dimensional image viewer, our application provides a seamless workflow for:

1. **AI-Assisted Annotation:** Integrating pre-trained models for automated segmentation using tiling strategies.
2. **Preprocessing & Refinement:** Interactive image enhancement and mask editing tools.
3. **Biological Quantification:** Calculation of density and transect-based metrics essential for botanical analysis.
4. **Batch Processing:** Automated inference and metric extraction across full folders of images.

2 Data Acquisition and Dataset

The development and validation of our framework relied on a specialised dataset of plant microscopy images. The biological material corresponds to leaves of the species *Euphorbia lathyris* (caper spurge), genotype SY03, provided by the *Instituto de Biología Molecular y Celular de Plantas* (IBMCP).

2.1 Sample Preparation and Imaging

A standardised protocol was followed to ensure consistent image quality. Leaves were collected from the third or fourth node from the plant's apex. The basal third of each leaf was selected, and the abaxial epidermis was carefully removed (peeled) to expose the internal tissue structures. Samples were decolourised overnight in a methanol-acetone mixture and subsequently stained with 0.05% Ruthenium Red for 30 minutes to highlight pectins in cell walls and laticifers.

Imaging was performed using a Leica MZ16F *MacroFluo* fluorescence stereomicroscope with a GFP filter. The raw

output consisted of high-resolution single-channel fluorescence images where laticifers appear as bright, interconnected structures against a darker background.

2.2 Dataset Characteristics

The complete dataset comprises 73 high-resolution images, each with dimensions of 3840×2400 pixels.

- **Labeled Subset:** 35 images were manually annotated by expert biologists. These ground truth masks delineate the full extent of the laticifer network and were used for supervised training of the model.
- **Unlabeled Subset:** An additional 38 images were provided without annotations.

Although the dataset contains 73 full-resolution leaves, patch extraction substantially increases the number of training samples. With a tile size of 512×512 and overlap, each image contributes many local patches, improving gradient diversity and helping the model learn fine laticifer patterns.

Several factors make this dataset particularly challenging for automated analysis: (i) extreme class imbalance, as laticifer pixels occupy only 4.6% of the total image area; (ii) structural complexity, with thin, branching networks of variable thickness requiring models to preserve topological connectivity; and (iii) label noise, as the manual annotation difficulty introduces inherent imperfections.

3 System Architecture

The system is designed as a modular desktop application using Python. The core architecture separates the User Interface (UI), the Deep Learning Inference Engine, and the Data/Quantification logic.

3.1 Frontend and Interaction Layer

The frontend is built with qtpy and integrated as a dock widget in napari, a fast, GPU-accelerated viewer for multidimensional images. This choice allows for the rendering of large microscopy images without performance degradation.

A central control widget serves as the primary hub, listening for layer insertion or removal to dynamically enable or disable tools based on the current context (e.g., quantification controls only activate once a mask layer exists).

3.2 Deep Learning Backend

The inference logic is decoupled from the UI to ensure responsiveness. We utilise PyTorch [Paszke *et al.*, 2019] for the model backend.

The core segmentation model is built upon the U-Net architecture [Ronneberger *et al.*, 2015]. To enhance feature extraction capabilities, we employ an SE-ResNeXt-50 encoder [Hu *et al.*, 2018]. This encoder is initialised with weights pre-trained on the ImageNet dataset [Deng *et al.*, 2009], facilitating faster convergence and improved generalisation despite the limited size of the botanical training dataset. Evaluated via 5-fold cross-validation, this configuration achieves a connectivity-aware Dice (clDice) [Shit *et al.*, 2021] of 0.58 and a standard Dice [Dice, 1945] of 0.49, ensuring robust topological preservation for downstream quantification.

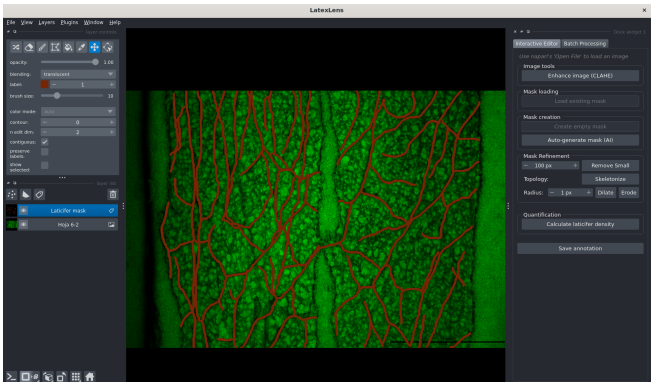

Figure 1: The LatexLens graphical user interface. The central panel displays the microscopy image with the AI-generated segmentation mask overlaid. The dock widget on the right provides controls for image enhancement, mask generation, refinement, and biological quantification settings.

Given the large dimensions of the input images, whole-image inference is computationally infeasible on standard GPUs. We implement a tiled inference strategy. The image is padded to ensure divisibility by the patch size (default 512×512). A sliding window with a defined stride (default 256) extracts patches. To mitigate edge artifacts common in tiled inference, we employ an overlapping strategy where predictions in overlapping regions are accumulated and averaged.

4 Annotation Workflow

The application supports a complete pipeline, illustrated in Figure 2, from image loading to the export of quantified data.

4.1 Preprocessing

Microscopy images often suffer from uneven illumination and low contrast. We integrated an image enhancement module directly into the UI. It utilises Contrast Limited Adaptive Histogram Equalisation (CLAHE) [Pizer *et al.*, 1987].

In the backend, the system automatically handles data type conversion and channel reduction (RGB to Grayscale) before applying CLAHE. For RGB data, the grayscale conversion uses the green channel, matching the fluorescence signal used in acquisition and avoiding channel mixing artefacts.

4.2 AI-Assisted Mask Generation

Users can trigger the "Auto-generate mask" function. This calls the backend inference engine described in Section 3.2. The resulting binary mask is injected into the viewer as a new interactive label layer. Crucially, the layer is set to 'editable' mode immediately. This supports a Human-in-the-Loop approach where the AI provides a "warm start" (often 80-90% accurate), and the expert biologist corrects false positives or connects broken laticifers using Napari's built-in painting tools. The main interface, displaying a segmented image and the control widget, is shown in Figure 1.

4.3 Batch Processing

For large experiments, the system provides a batch-processing tab that accepts an input directory of microscopy images and outputs a mirrored directory of predicted masks. The batch pipeline performs tiled inference for each image and records per-image metrics into a CSV file. This feature enables researchers to process entire datasets without manual interaction while still preserving a structured audit trail.

4.4 Quantification Suite

A key contribution of this work is the integration of domain-specific biological metrics. Raw binary masks are insufficient for biological analysis; derived metrics are required. We implement two primary quantification methods:

Pixel Density Ratio

This metric calculates the area fraction occupied by laticifers:

$$D_{pixel} = \frac{\sum_{x,y} \mathbb{I}(M_{x,y} > 0)}{H \times W} \times 100 \quad (1)$$

where M is the segmentation mask. To avoid background bias when images contain surrounding black borders, the tool can automatically derive a tissue-only region from the laticifer network using a convex-hull [Graham, 1972] or morphological envelope, effectively cropping the measurement to the leaf area before computing density.

Transect Intersection Method

In stereology and plant anatomy, network density is often estimated by counting intersections with test lines (transects). We digitise this method within the quantification module. The user specifies the number of lines (N) and direction (horizontal/vertical/both). The system generates equidistant lines across the image. For each line L_i, we analyse the pixel intensity profile. An intersection is detected as a transition from background (0) to foreground (1).

$$Intersections(L_i) = \sum_j [P_j = 0 \wedge P_{j+1} = 1] \quad (2)$$

In this expression, L_i is the i-th transect line, P_j is the binary mask value sampled at position j along that transect, and the logical condition $[P_j = 0 \wedge P_{j+1} = 1]$ counts each background-to-laticifer transition as one intersection.

The application visualises these transects as vector shapes overlaid on the image and highlights intersection points. The statistics (mean intersections per line $\pm$ std) are reported to the user and logged to the CSV.

5 Implementation Details

5.1 Handling High-Resolution Imagery

The GUI ensures thread safety when interacting with the viewer. To prevent memory overflows during inference on 4K resolution images, the sliding window approach processes tensors on the GPU and offloads accumulated results to CPU memory where necessary.

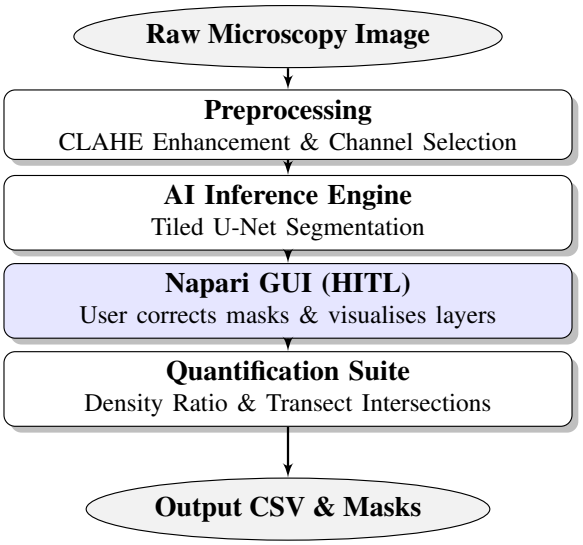

Figure 2: The Annotation Workflow. The system processes raw images through enhancement and AI segmentation before presenting them in the Napari interface. Users can interactively refine the masks (Human-in-the-Loop) before automated quantification.

5.2 Data Persistence

A dedicated persistence routine ensures data integrity and creates a robust audit trail for biological experiments. When the user presses the saving button, the system performs the following operations atomically:

1. Saves the original state of the image.
2. Saves the current state of the mask (potentially manually refined).
3. Recalculates all metrics (Density and Transect) based on the original image and the saved state of the mask.
4. Appends a row to a CSV log file with timestamps, dimensions, and computed metric.

6 Discussion and Future Work

We presented LatexLens, a tool that democratizes access to deep learning segmentation for plant biologists. By wrapping complex inference logic in a familiar GUI and integrating standard biological quantification metrics, we significantly reduce the time required for laticifer analysis.

The modular design allows for future extensions. Currently, the system uses a 2D U-Net. Future work includes integrating 3D segmentation capabilities for confocal microscopy stacks.

Acknowledgments

This work has been supported by CIPROM/2022/6 (FASSLOW), funded by Generalitat Valenciana, and Spanish grant PID2024-162030OB-100 (ROBIN), funded by MCIN/AEI/10.13039/501100011033 and ERDF "A way of making Europe", Cátedra ENIA-UPV in Sustainable AI Development, TSI-100930-2023-9.

References

[Agrawal and Konno, 2009] Anurag Agrawal and Kotaro Konno. Latex: A model for understanding mechanisms, ecology, and evolution of plant defense against herbivory. *Annu. Rev. Ecol. Evol. Syst*, 40:311–31, 12 2009.

[Chiu *et al.*, 2022] Chi-Li Chiu, Nathan Clack, and the napari community. napari: a python multi-dimensional image viewer platform for the research community. *Microscopy and Microanalysis*, 28(S1):1576–1577, 08 2022.

[Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

[Dice, 1945] Lee R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.

[Fahn, 1979] Abraham Fahn. Secretory tissues in plants. 1979.

[Graham, 1972] Ronald L. Graham. An efficient algorithm for determining the convex hull of a finite planar set. *Inf. Process. Lett.*, 1:132–133, 1972.

[Hagel *et al.*, 2008] Jillian Hagel, Edward Yeung, and Peter Facchini. Got milk? the secret life of laticifers. *Trends in plant science*, 13:631–9, 11 2008.

[Hu *et al.*, 2018] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.

[LeCun *et al.*, 1989] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.

[Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

[Pizer *et al.*, 1987] Stephen M. Pizer, E. Philip Amburn, John D. Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B. Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3):355–368, 1987.

[Ronneberger *et al.*, 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing.

[Shit *et al.*, 2021] Suprosanna Shit, Johannes C. Paetzold, Anjany Sekuboyina, Ivan Ezhov, Alexander Unger, Andrey Zhylka, Josien P.W. Pluim, Ulrich Bauer, and Bjoern H. Menze. clDice: A Novel Topology-Preserving Loss Function for Tubular Structure Segmentation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16560–16569, 2021.