

Splanting: 3D plant capture with gaussian splatting

TOMMY OJO, University of Saskatchewan, Canada

THAI LA, University of Saskatchewan, Canada

ANDREW MORTON, University of Saskatchewan, Canada

IAN STAVNESS, University of Saskatchewan, Canada



Fig. 1. 3D gaussian splatting based capture of a canola plant in a laboratory environment with a visually cluttered background: (left-to-right) reconstruction including background splats (scrambled for privacy), isolated canola *splant* after removal of background splats by spatial clustering, isolated flower clusters after radiance-based segmentation, and flower clusters in grid.

3D capture and characterization of plant shoot architecture is a grand challenge in plant phenotyping research, made difficult by plants' intricate 3D shape, composed of thin and flat sub-structures (stems, leaves, flowers, pods, etc.). In this paper, we show that 3D gaussian splatting is well-suited for capturing 3D plant representations, which we call *splants*. We report a simple and fast capture procedure and 3DGS processing software that is tailored to foreground object capture. Splant generation worked well across plant species and growth stages. Our preliminary results point to a promising future for plant phenotyping, which we expect will lead to a dramatic increase in the use of multi-view imaging and 3D analysis in plant pathology and plant breeding research.

CCS Concepts: • **Computing methodologies** → **Point-based models; Reconstruction**; • **Applied computing** → **Agriculture**.

Additional Key Words and Phrases: 3D gaussian splatting, radiance field rendering, image processing, 3D reconstruction, plant shoot architecture, plant phenotyping

1 Introduction

3D reconstruction and rendering of a plant's above-ground shoot architecture is a long-standing grand challenge in plant phenotyping. The shape and spatial arrangement of a plant's shoots determine its capacity to capture solar energy, its yield potential, its susceptibility to certain plant pathogens, and how easily it can be harvested. However, characterizing a plant's shoot architecture is notoriously difficult due to its intricate 3D shape, composed of lower-dimensional manifolds: 1D branches, stems, silique and 2D leaves, flower petals, pods. Reconstruction of plants with traditional multi-view photogrammetry is made difficult by self-occlusion and self-similarity of plant structures.

Authors' Contact Information: Tommy Ojo, tommy.ojo@usask.ca, University of Saskatchewan, Saskatoon, SK, Canada; Thai La, thai.la@usask.ca, University of Saskatchewan, Saskatoon, SK, Canada; Andrew Morton, andrew.morton@usask.ca, University of Saskatchewan, Saskatoon, SK, Canada; Ian Stavness, ian.stavness@usask.ca, University of Saskatchewan, Saskatoon, SK, Canada.

3D gaussian splatting (3DGS) [Kerbl et al. 2023] has generated impressive visual results for natural scenes, but has yet to be evaluated for complex plant architectures with substantial self-occlusion. In this paper, we provide a preliminary evaluation and analysis of 3DGS-based plant reconstructions, which we call *splants*. We examine *splants* across a range of plant species and demonstrate that their optimized splats are representative of physical plant structures and that splat location, orientation and radiance can be used to localize and identify plant organs.

The contributions of our paper include: 1) an exemplar *splant* dataset including diverse samples across plant species and growth stages; 2) a simple *splant* capture procedure; 3) *splant* software for 3DGS processing (optimized for foreground object capture) and splat filtering, clustering, and contraction processing, towards skeletonization of a plant’s architecture; 5) a demonstration of using splat features (location, orientation, radiance) for plant organ segmentation; and 6) a discussion of the implications of *splant* capture for plant phenotyping.

2 Related Work

3D Plant Capture. Prior to recent advances in radiance field rendering, a wide range of approaches have been proposed for 3D capture of plant shoot architecture. Plants have been subjected to multi-view photogrammetry [Drofova et al. 2023], X-ray imaging [Dutagaci et al. 2020], and laser scanning [Schunck et al. 2021]. Of these traditional capture methods, 3D laser scanning has worked best [Young et al. 2024], but still requires fairly slow and careful manual or robot-controlled movement of the scanner to capture the thin, intricate, and self-occluded structures that comprise a plant’s overall architecture.

Neural Radiance Fields for Plants. Radiance field rendering research is rapidly accelerating in many directions and quickly gaining interest in specialized domains, such as plant phenotyping. Neural Radiance Field (NeRF) capture of plants has shown promising initial results for specific species [Saeed et al. 2023] and even in challenging agricultural field conditions [Arshad et al. 2024].

3D Gaussian Splatting for Plants. 3DGS has been successfully used to measure stem length and count bolls in cotton plants [Jiang et al. 2024]. However, to the best of our knowledge, 3D gaussian splatting has yet to be evaluated for complex plant structures that include substantial self-occlusion and for multiple plant species. Concurrent to our work, 3DGS and NeRF reconstructions of wheat plants have been compared [Stuart et al. 2024]. One advantage of 3DGS compared to implicit representations, like NeRF, is that it directly generates an explicit representation of the plant as a 3D splat cloud that is suitable for downstream 3D plant simulation [Cieslak et al. 2022; Huang and Stavness 2022; Li et al. 2023].

3 Methods

We have designed a quick and effective workflow for *splant* generation that extends the canonical 3D gaussian splatting optimization procedure for capturing individual plants from multi-view images.

3.1 Splant Capture

Imaging. Our *splant* capture protocol uses a single 360° side-view video clip around the plant (Figure 2), from which we extract ~100 frames for 3DGS processing. Our pilot testing with different cluttered backgrounds, lighting conditions, and camera paths all worked well for static scenes. We found the traditional setup of a fixed camera and a rotating plant on a turn-table, as well as entirely blank backgrounds, were less effective because they degraded the initial structure-from-motion (SfM) step of the 3DGS process. As future work, fiducial markers or a vection background may improve SfM initialization.

3DGS Processing. We use the original implementation of 3DGS for processing *splants* [Kerbl et al. 2023]. Although our interest is only in the foreground plant, background features assist in the SfM estimation of camera parameters and initial splat locations. However, to reduce processing time and focus 3DGS processing on the plant




Fig. 2. Our *splant* capture protocol uses dense side-view images or a video clip around the plant.

Table 1. Exemplars from our *splant* dataset.

species	growth stage	# images in	# splats out	shown
canola	rosette	77	48,264	Fig. 5
canola	flowering	97	214,598	Fig. 1
wheat	mature	90	104,433	Fig. 3
bean	vegetative	93	144,486	supp. video
chickpea	vegetative	82	150,080	Fig. 3
lentil	vegetative	83	431,979	Fig. 4
pea	vegetative	103	127,031	Fig. 7
pea	Pods	79	187,807	Fig. 7
pea	mature	93	143,775	Fig. 7

subject, we pre-process input images by cropping to the plant region-of-interest (ROI) using a COCO-pretrained YOLOv8n detector [Jocher et al. 2023]. For the pea splant, ROI image cropping reduced 3DGS processing time by approx. 25% and reduced the number of background splats by half (Table 2) although this was scene dependent. Tight ROI cropping, however did cause SfM failures in some cases, in which case we revert to the original images.

Background splat removal. The primary post-processing step is to isolate the foreground plant splats from the background splats and to remove spurious splats that are not associated with plant subjects. Using the CloudCompare v2.11.3 CLI (CC), we first perform Statistical Outlier Removal (SOR) to remove the “floating” splats using a standard deviation multiplier threshold of 1 and 15 points for mean distance estimate. Next, we generate candidate splat clusters via connected components in CC, with one component representing the plant and the other components representing the background. We manually select the component that best represents the plant in CC, a step which we plan to automate in future work. After exporting the selected plant component, we merge back per-splat parameters, which is not yet supported in CC. Of note, the support surface for the plant was often included in the plant component, which generated rather pleasing renderings of wire shelves and carts (Figures 1a, 4), but had to be manually removed.



Fig. 3. Example renderings of wheat and pea splants, showing side, top-down, and zoomed-in views.

3.2 Splant Dataset

We have captured splants across a range of agronomically-important plant species, growth-stages, and backgrounds. While the present dataset is modest in size, we expect it to grow rapidly due to the ease and speed of capture, supported by our splant processing software. Exemplar samples from the dataset are summarized in Table 1.

3.3 Splant Analyses

We have developed preliminary analyses of splant renderings to demonstrate their potential utility for characterizing plant traits. In order to demonstrate that the optimized splats that comprise a splant are representative of plant structures, we explore the use of splat locations to skeletonize splants and compute splant volume, splat radiance to isolate canola flowers, and splant orientation to differentiate splat clusters belonging to different types of organs.

4 Results

Qualitative results. Figure 1 depicts a representative rendering of a canola splant at the flowering stage, first with the whole room background reconstructed, then with plant-related splats isolated, and finally with canola flower splats isolated based on their radiance. Additional exemplars include wheat and pea splants with visible spikes and pods (Figures 3), lentil and chickpea splants with small dense leaflets (Figure 4), and a pea splant over time (Figure 7).

Using splat location to skeletonize plant architecture. Using splant location alone is akin to point-cloud processing [Meyer et al. 2023]. Splant splat clouds are substantially sparser than traditional plant point clouds, but utilizing similar contraction-based processing of splant locations clearly reveals the stems and leaves in a rosette-stage canola splant with only a few missing regions due to self-occlusion (Figure 5). We also found that the concave hull of splat locations provided a tight fit for the chickpea splant and may be a good proxy for a plant’s above ground biomass volume.



Fig. 4. Example renderings of lentil and chickpea splints, showing side, top, and zoomed-in views. The small leaflets in these species are visible, but blurry when zoomed-in.




Fig. 5. A rosette-stage canola splint: (a) oblique perspective, (b) top-down perspective, (c) splat locations after contraction showing that splats are arranged near plant structures, (d) stem splats isolated using RANSAC.

Using splat orientations to differentiate plant organ types. When optimized to render thin and flat plant structures (stems and leaves), we expect that splats will be oriented to align with thin structures. To characterize the optimized splint splats, we measure the distribution of their orientations for three canonical organ/objects within our splint

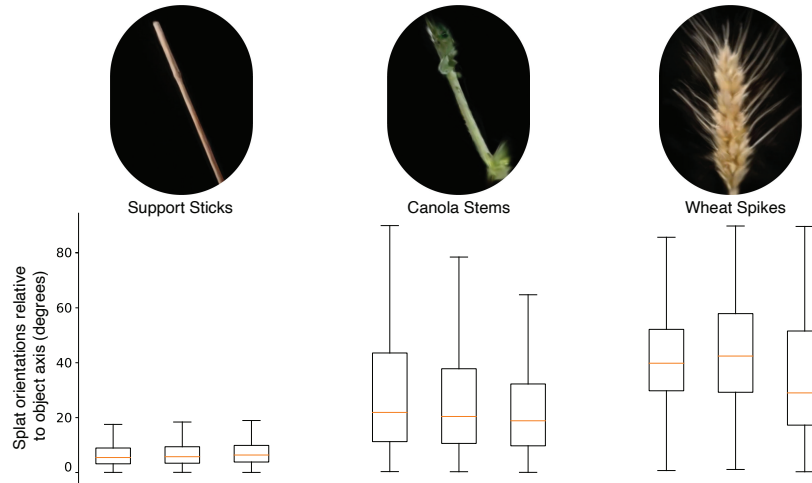


Fig. 6. Orientation of splat principle axis relative to the overall object axis for three instances of artificial support sticks, canola stems, and wheat spikes (isolated from full plant captures). Splats are aligned to thin plant organs, and uniformly oriented for wheat spikes to capture the thin awns.

Table 2. Comparison of pea splat generation with and without cropping the input images to the plant ROI.

input images	# splats total	# splats background	# splats plant	3DGS processing time (minutes)
original	515,971	388,940	127,031	29
cropped	372,541	114,658	257,883	22

dataset: wheat spikes, canola stems, and artificial support sticks. We first find the principal axis of the organ using principal component analysis on the isolated splat cloud. Each individual splat geometry is composed of its scaling factors and the rotation matrix created from the normalized given quaternions. Using this information, we can efficiently extract the principal axis of each splat by extracting a column from the rotation matrix that corresponds to the largest scaling factor. We found that splats were well aligned to the stick, mostly aligned to the stem, and more uniformly distributed for the wheat spike (Figure 6).

Using splat radiance to isolate canola flowers. Given their distinctive yellow color, canola flowers are the “low-hanging fruit” of plant phenotyping. However their dense arrangement into cylindrical clusters (racemes), makes it difficult to accurately count individual flowers from a single image. We isolate flowers in the canola splat by extracting the splats’ 0th spherical harmonic band (base color) and thresholding for yellowness in the LAB colorspace. Additional filtering was done for low-opacity splats and by SOR to remove small parts of stems that appear yellow. We show isolated canola flowers in place and zoomed-in for nine examples (Figure 1).

Number of input views required. We found that our side-view 360° capture procedure required a narrow separation between input images to work well. For most splats, less than ~30 images (12° separation) caused the initial SfM step to fail, and greater than ~100 images (~3.5° separation) did not improve the reconstructed splat.



Fig. 7. The same pea plant captured across growth stages: vegetative, pod-development, senescence.

5 Discussion and Implications

To conclude, we discuss the implications of *splants* for plant phenotyping research. Our preliminary attempts at generating 3D splants resulted in highly realistic 3D renderings for every species attempted, with few failures. Compared to laser scanning [Schunck et al. 2021], splant capture is faster (~15 sec vs. ~15 min), cheaper (\$100 camera vs. \$10,000+ scanner), generates smaller point clouds (~100k splats vs. ~1M+ points), and is less tedious. In our exemplar splants, we have observed some limitations: splats were missing in highly occluded regions, support structures (cart, tabletop) were not always removed with our simple SOR approach, and very dense plants did not admit a view of internal architecture and rendered somewhat fuzzy (Figure 4). However, by and large the splant exemplars we produced are rather wonderful at a whole plant level, and have sufficient splat density and detail to permit interrogation of whole plant architecture and individual plant organs.

The implications of splant capture for plant phenotyping are numerous. The speed and ease of capture would permit 3D phenotyping for modest size plant populations (100s-1000s of individuals) for plant breeding and pathology experiments. The computational resources required (a desktop GPU) are accessible to most plant science research groups around the world. The resulting photorealistic renders are likely sufficient for qualitative visual phenotyping and screening, which is common and often preferred by breeders. High-throughput and automated splant capture and analysis will require adapting phenotypic analysis approaches from 2D images and 3D point clouds to 3D splat clouds, but the compact representation afforded by splants will make analysis workflows more scalable. 360° splant capture is only feasible for certain plant experiment conditions and growth facilities, but these types of experiments, e.g. plant pathology, are the ones that require detailed phenotypic information potentially afforded by splants. Evaluating outdoor splant capture for crops grown in dense field plots is an important future step for expanding splant use in crop breeding programs. Finally, new 3DGS variants are rapidly improving efficiency and accuracy of reconstructions for general scenes, and optimizing these methods for *splanting* is a promising direction for future work.

Acknowledgments

We thank collaborators who provided plants, including Kirstin Bett, Adam Carter, Devan Guenter, Randy Kutcher, Isobel Parkin, Steve Shirtliffe, Ana Vargus, Lipu Wang, and Haiying Yuan as well as Isaac McKay, Amanda Ewen, Rodrigo Godoy, Dawn Omoluabi, Matthew Drotor, and Lingling Jin. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada.

References

- Muhammad Arbab Arshad, Talukder Jubery, James Afful, Anushrut Jignasu, Aditya Balu, Baskar Ganapathysubramanian, Soumik Sarkar, and Adarsh Krishnamurthy. 2024. Evaluating neural radiance fields (NeRFs) for 3D plant geometry reconstruction in field conditions. *Plant Phenomics* in press (2024). <https://doi.org/10.34133/plantphenomics.0235>
- Mikolaj Cieslak, Nazifa Khan, Pascal Ferraro, Raju Soolanayakanahally, Stephen J Robinson, Isobel Parkin, Ian McQuillan, and Przemyslaw Prusinkiewicz. 2022. L-system models for image-based phenomics: case studies of maize and canola. *in silico Plants* 4, 1, Article diab039 (2022), 19 pages.
- Irena Drofova, Haozhou Wang, Wei Guo, Martin Pospisilik, Milan Adamek, and Jan Valouch. 2023. 3D reconstruction of a group of plants by the ground multi-image photogrammetry method. In *33rd International Conference Radioelektronika (RADIOELEKTRONIKA)*. IEEE, 1–4.
- Helin Dutagaci, Pejman Rasti, Gilles Galopin, and David Rousseau. 2020. ROSE-X: an annotated data set for evaluation of 3D plant organ segmentation methods. *Plant Methods* 16 (2020), 1–14.
- Danny Huang and Ian Stavness. 2022. Large Growth Deformations of Thin Tissue Using Solid-Shells. *IEEE Transactions on Visualization and Computer Graphics* 29, 3 (2022), 1893–1909.
- Lizhi Jiang, Changying Li, Jin Sun, Peng Chee, and Longsheng Fu. 2024. Estimation of cotton boll number and main stem length based on 3D gaussian splatting. In *2024 American Society of Agricultural and Biological Engineers Annual Int. Meeting*. 1.
- Glenn Jocher, Ayush Chaurasia, and Jing Qiu. 2023. *Ultralytics YOLO*. <https://github.com/ultralytics/ultralytics>
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.* 42, 4, Article 139 (2023), 14 pages.
- Bosheng Li, Jonathan Klein, Dominik L Michels, Bedrich Benes, Sören Pirk, and Wojtek Palubicki. 2023. Rhizomorph: The coordinated function of shoots and roots. *ACM Trans. Graph.* 42, 4, Article 59 (2023), 16 pages.
- Lukas Meyer, Andreas Gilson, Oliver Scholz, and Marc Stamminger. 2023. CherryPicker: Semantic skeletonization and topological reconstruction of cherry trees. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6244–6253.
- Farah Saeed, Jin Sun, Peggy Ozias-Akins, Ye Juliet Chu, and Changying Charlie Li. 2023. PeanutNeRF: 3D radiance field for peanuts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6254–6263.
- David Schunck, Federico Magistri, Radu Alexandru Rosu, André Cornelißen, Nived Chebrolu, Stefan Paulus, Jens Léon, Sven Behnke, Cyrill Stachniss, Heiner Kuhlmann, et al. 2021. Pheno4D: A spatio-temporal dataset of maize and tomato plant point clouds for phenotyping and advanced plant analysis. *PLOS ONE* 16, 8, Article e0256340 (2021), 18 pages.
- Lewis AG Stuart, Darren M Wells, Jonathan A Atkinson, Simon Castle-Green, Jack Walker, and Michael P Pound. 2024. High-fidelity Wheat Plant Reconstruction using 3D Gaussian Splatting and Neural Radiance Fields. *GigaScience* pre-print (2024), 1–14.
- Therin J Young, Shivani Chiranjeevi, Dinakaran Elango, Soumik Sarkar, Asheesh K Singh, Arti Singh, Baskar Ganapathysubramanian, and Talukder Z Jubery. 2024. Soybean canopy stress classification using 3D point cloud data. *Agronomy* 14, 6 (2024), 1181.