

Introduction

We have developed an open-source automatic pipeline called AutoSplant that turns cellphone video into 3D models of plants. These models are 3D Gaussian Splats [3], which are quick to scan and render while preserving the 3D structure. Our aim is to improve the accessibility of 3DGS technology.

What is 3DGS?

3D Gaussian Splatting (3DGS) [3] is a technique that uses backpropagation to train a 3D model. This model is stored as a collection of thousands/millions of individual Gaussians instead of traditional textures and meshes.



Three individual Gaussians. They have soft borders because the colour intensity follows a normal distribution.

Executive Summary

Problem:

- Current techniques for manipulating 3DGS require a lot of labour. Newer techniques exist but are less accessible.
- Robust and open-source processing pipelines do not exist and would greatly increase scalability.

Goals:

- Make 3DGS end-to-end software that is open and accessible
- Improve the accessibility of the technology and automate the process.

Key Contribution:

- We have developed an open source 3DGS processing pipeline that turns simple camera video into segmented 3DGS models.
- We use YOLO [4] image masks in conjunction with a depth map to provide SAGA with point samples on the 3DGS model for segmentation.

Key Results:


- The technology is potent; it has a high success rate but can fail under certain circumstances.
- We have identified many areas for further improvements.

Related Works

AutoSplant is similar to consumer pipelines, like XGrids' LCC Studio, which provide easier ways to generate/manipulate 3DGS. AutoSplant is a bit different:

- AutoSplant includes automatic/assisted segmentation in its pipeline
- AutoSplant uses cellphone video / images instead of an expensive or proprietary scanner.
- Since the pipeline/device is not proprietary, customization like colour correction can be performed.
- AutoSplant is free, open-source, and it runs locally, keeping your data private.

Methodology



COLMAP Step:

- AutoSplant takes a video or a collection of multi-view images as input.
- Finds the camera position of each image: A coordinate in space and rotation.
- Generates a point cloud of the scene.
- Aligns the orientation so that the scene, and thus the final object, is level with the ground.

What is COLMAP?

COLMAP is an open-source, general purpose Structure-from-Motion (SfM) pipeline. [5]

It is used frequently to make 3D reconstructions from images and video.



PGSR Step:

- The scene reconstruction and images are used to train a 3DGS Model using PGSR.
- The lighting, colours, and other visual details are baked in during this step. Modifying 3DGS Models can be difficult and labour intensive.
- Gaussians use Spherical Harmonics to look different from different angles.
- This image is not a photo; it is a 3DGS Model!

What is PGSR?

Planar-Based Gaussian Splatting (PGSR) is an improved technique for generating 3DGS Models [2].

It confines the individual Gaussians to detected surfaces, accurately preserving the structure of the model.



SAGA Step:

- AutoSplant uses SAGA to give every Gaussian an affinity feature [1].
- Gaussians with similar affinity features belong to the same object [1].
- Affinity features have a 'scale' dimension, which allows a single Gaussian to belong to multiple objects [1]. Scale is given with the sample point to begin segmentation.

What is SAGA?

Segment Any 3D Gaussian (SAGA) is a technique for segmenting 3DGS Models [1].

It identifies every object in the input video, and 'retrains' the model to include an affinity feature for each individual splat.



AutoSplant Step:

- YOLO (version 26L-seg) identifies objects and generates a mask [4], which is an outline of the objects.
- A depth map is used to find the holes in the mask.
- This step provides a list of sample points on the object from which to begin the SAGA segmentation [1].

What is YOLO?

You Only Look Once (YOLO) is a 2D object classification model [4].



Evaluation & Results

A 1-2 minute cellphone video of an object is often good enough, as long as many angles were captured. We found the best scanning path to be a spiral around an object, in order to capture multiple heights. Visit biglab.usask.ca/gallery/splants to see

The quality of the result heavily depends on the quality of the input data. Blurring, moving objects, or a bad scanning path leads to artifacts or a complete failure.



Bird of Paradise




Bean



Dragon Tail



Chickpea



Wheat

Discussion

AutoSplant is open-source and runs locally:

- Proprietary products can act as surveillance devices, sending sensitive data back to 3rd party servers.

AutoSplant inherits many of the limitations of its parts:

- PGSR has issues with low-detail surfaces, and sometimes generates artifacts anyways [2].
- SAGA and YOLO can have noisy segmentations, due to inaccurate 2D segmentation [1][4].

Many other improvements could also be made:

- Li-DAR data could be incorporated into the COLMAP step, allowing objects to be to-scale with their real counterparts [5].

Multi-Object segmentation is possible:

- It should be possible to get every plant in a 3D scene to be automatically segmented, opening up applications to greenhouse scanning.



3DGS Model of a colour corrected lysimeter greenhouse

References

1. Cen, J., [...] & Tian, Q. (2025). *Segment Any 3D Gaussians*.
2. Chen, D., [...] & Zhang, G. (2025). PGSR: Planar-Based Gaussian Splatting [...] *IEEE Transactions on [...] Computer Graphics*, 31(9), 6100–6111.
3. Kerbl, B., [...] & Drettakis, G. (2023, August 8). *3D Gaussian Splatting for Real-Time Radiance Field Rendering*.
4. Redmon, J., [...] & Farhadi, A. (2015, June 8). *You Only Look Once: Unified, Real-Time Object Detection*.
5. Schonberger, J. L., & Frahm, J.-M. (2016). Structure-from-Motion Revisited. *2016 IEEE Conf on Computer Vision and Pattern Recognition (CVPR)*, 4104–4113.

Acknowledgements

A special thanks to the following for discussing the potential uses and providing access to their crops/plants/experiments for scanning:

Steve Shirtliffe, Adam Carter, Jamie Otterson, Daniela Moraes, Ali Lotfi, Sandesh Neupane, Kirstin Bett, Ana Vargus, Haiying Yuan, Devan Guenter, Sally Vail, Steve Robinson, and Isobel Parkin.

AI Disclosure:

Generative AI or Large Language Models (LLMs) were **not** used at any point in this project. All ideation, coding, graphic design, or otherwise was done by a human.