Author:
Shashwat Sahay

Supervisors:
Petr Kellnhofer and Lukas Uzolas

Responsible Professor:
Elmar Eisemann

TU Delft

ISOLATING A TREE'S SKELETON USING A 3-DIMENSIONAL RECONSTRUCTION

Introduction

L-Systems are grammars that allow for the procedural generation of trees built to look a certain way. To be able to derive an L-grammar from a real tree, we first need to understand how that tree is structured. The idea behind my research is to develop a pipeline that finds this structure, using a 3-dimensional respresentation of the tree, which gives us more information to work with as compared to only 2 dimensions

NeRF

To recreate the tree in 3 dimensions, I used NeRF, a methodology that uses neural networks to be able to output renderings of a scene given an input viewing location and direction, after being trained on multiple images with known camera positions. We can then integrate over a large number of "cameras" to derive a point cloud of the scene (or in this case, part of it)

Filtering and Skeletonization

After producing the point cloud, to reveal the structure of the tree, we have to be able to remove its leaves. In my paper, I settled on a simple color-based filtering, but I also outlined multiple other methods that could be used to improve on the pipeline as it is.

To then skeletonize the tree, I used an algorithm called Laplacian-Based Contraction, which "contracts" the points in the cloud closer to each other until they have 0 volume, and then finding the topology of this contracted point cloud using a minimum spanning tree algorithm. This approach was also compared to a medial axis approach in the paper

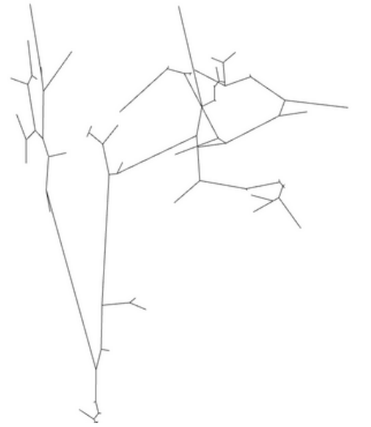

Evaluation

Here I've overlayed the topology of the tree with the tree itself, as well as a ground truth that I annotated by hand. We can qualitatively see that it does cover almost all the branches of the tree, but then continues further into the leaves, which means that the filtering step has great room for improvement.