

Generating real-time detailed ground visualisations from sparse aerial point clouds

Aidan Murray, Scarlet Mitchell, Alexander Bradley
www.cobrasimulation.com

Eddie Waite, Caleb Ross, Joanna Jamroz, Kenny Mitchell
www.napier.ac.uk

Cobra Simulation Ltd,
Bathgate

Edinburgh Napier University and Cobra Simulation Ltd,
Edinburgh




Figure 1: Comparison of similar walking eye level viewpoints, between Google Earth (left), the Epic Games' Unreal Engine 5.0 LiDAR plugin illustrating source point data (middle) and our method (right).




Figure 2: Tree reconstruction inferred from sparse point cloud data showing before and after addition of noise-based SDF natural shape detail.




Figure 3: The combination of extended point surface star-like capsules forms resulting tree structures inferred from sparse 1m resolution data.

1 Procedural LiDAR point cloud detail amplification

Building realistic wide scale outdoor 3D content with sufficient visual quality to observe at walking eye level or from driven vehicles is often carried out by large teams of artists skilled in modelling, texturing, material shading and lighting, which typically leads to both prohibitive costs and reduced accuracy honoring the variety of real world ground truth landscapes. In our proposed method, we define a process to automatically amplify real-world scanned data and render real-time in animated 3D to explore at close range with high quality for training, simulation, video game and visualisation applications.

We perform procedural 3D world generation processing upon remote sensed aerial LiDAR scan data [2], which is often noisy, sparse and incomplete. However, given procedural, categorised shape function interpolation with signed distance fields (SDF) ray marching, the sparse data holds sufficient key information to reconstruct large areas of terrain rapidly formed within spanned volume extents, that encode nearest adjacent points, classification and material properties. That is, we infer rich local surface material, and shape details from simple point cloud data clusters by mapping matched 'templates' to high-quality pre-scanned local surfaces. Each template holds procedural variation so no two surfaces are exactly alike, and blend seamlessly together with noise and prior based distance field interpolation methods to combine and form a complete digital twin of the real world. We compare with Google Earth in figure 1.

A nearest relative of our work is GANcraft which takes the low fidelity blocks world representation of the video game Minecraft and generates towards realistic results with a neural rendering framework [1]. However, whilst our results are free of cubist style artifacts, the blocky outline appearance is not addressed by GANcraft.

The scheme takes points of spatial scene data and generates arbitrary 3D views with high density inferred shape corresponding to the original reference scanned location through point based distance field interpolated rendering, including surface color (albedo) from registered aerial maps.

Our lean new local point based extended volume element rendering descriptor, includes cached point adjacency of neighbouring points, along wide material and shading data intended for efficiently combining on GPU for real-time rendered reconstruction of scenery.

The signed distance field (SDF) based detailed shape definition, whose volume comprises the set of adjacency vectors to neighbouring point cloud points forming, e.g., capsules emanating from a central 3D location, and

may be any shape derivable from 'star'-like structured volume elements.

Detailed appearance reconstruction is achieved through authored analytic SDF shaping and noise functions according to common feature classifications, e.g. trees, grass, roads, etc. recovering spatially varying material appearance properties including diffuse, specular, roughness, etc.

Hierarchical spatial partitioning (3D grid buckets) is performed for compute shader efficient multi-threaded data transformation of the raw sparse point cloud data into efficient fragment shader rendering packets with up to eight neighbouring adjacent point cloud data.

We provide an efficient ray marched screen aligned quad rendering of the inferred detail SDFs with temporal re-projection depth based occlusion culling optimization, bounding sphere limit culling, and per chunk culling. Critically, to counter overdraw of naive quad rendering we perform occlusion culling with re-projection of approximate scene depth from the previous frame to sparingly cull at the level of quads.

Ground level natural surfaces are achieved through offsetting approximately shaped surfaces based on increasing frequencies of directed shape function according to level of detail, e.g., a vertically tapered noise perturbed grass surface. For trees and bushes, the extended point derived surface is generated with shape steered by the connectivity information of adjacent point cloud data of same foliage category (fig. 2). This adjacency data in our example forms oriented SDF capsules that, when composited, resolve to a detailed appearance of natural landscape features, e.g leaf-like features by offsetting the shape's surface in all directions (fig. 3).

Our new scheme provides a both texture-less and mesh-less framework, which affords memory efficient, flexible level of detail and visibility optimizations for rendering scenery in real-time. Generating rapid walk-through environments from efficiently gathered sparse LiDAR data forms a new avenue for digital twins for simulation and training, and entertainment, and provides a seed for future analytic and neural extended point-based rendering approaches.

- [1] Zekun Hao, Arun Mallya, Serge J. Belongie, and Ming-Yu Liu. GANcraft: Unsupervised 3d neural rendering of minecraft worlds. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14052–14062, 2021.
- [2] Nina Varney, Vijayan K Asari, and Quinn Graehling. Dales: A large-scale aerial lidar data set for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 186–187, 2020.