GPU driven Dense Reconstruction for Community Photo Collections

Jan-Michael Frahm, Pierre Fite-Georgel, David Gallup, Tim Johnson, Rahul Raguram, Changchang Wu,
Yi-Hung Jen, Enrique Dunn, Brian Clipp, Svetlana Lazebnik, Marc Pollefeys

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

ETH

Overview

- Efficient image grouping
- Geotag-driven scene grouping
- High performance geometric image registration
- Novel stereo and depth fusion

2.8M images (650 GB)

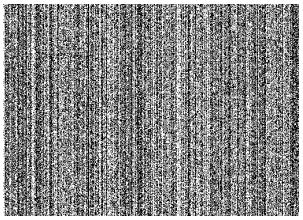

Compressed Representation
512-bit binary codes (176 MB)

Gist feature vectors are first computed by performing parallel image convolutions with CUDA. They are then compressed to binary codes using the CUBLAS library.

A CUDA implementation of the k-medoids algorithm provides an initial partition of the dataset. The resulting clusters are tested for geometric consistency. SIFT features are extracted using CUDA, and then a CUDA implementation of the RANSAC algorithm estimates the fundamental matrix between image pairs. Clusters with enough consistent images are considered "geometrically verified.". Each cluster is represented by an iconic image.

100K Clusters

Geometric Verification

~21K Clusters

At this point, many clusters may describe the same scene. Using GPS information we find iconics that are close to each other. Each of these connections is verified using SIFT and RANSAC as in the geometric verification phase. Iconics that describe the same scene are grouped together.

Geotagging Information
(10% of images)

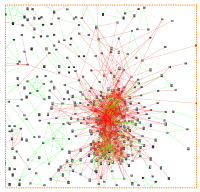

Clusters Connection

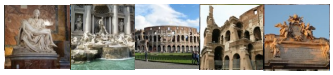

Consistent Set of Clusters

Sparse reconstructions are computed using an incremental structure from motion approach. Using the obtained calibrated view, stereo depthmaps are computed using a real-time GPU-accelerated multi-view planesweep method. Finally the depthmaps are fused using a multi-layer height-map model.

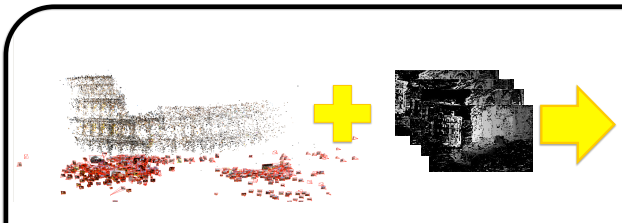

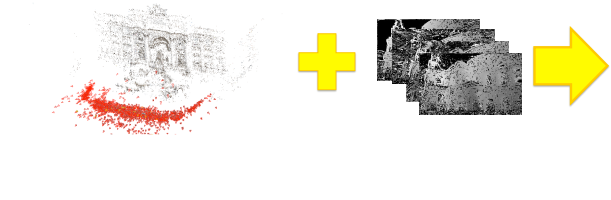

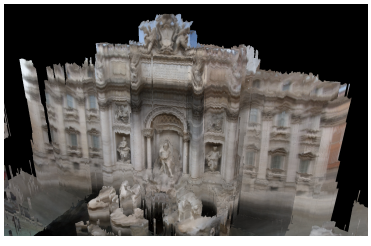

Sparse Reconstruction **Stereo Depthmaps** **Dense 3D Model**