Carnegie Mellon

MAP Visibility Estimation for Large-Scale Dynamic 3D Reconstruction

Hanbyul Joo, Hyun Soo Park, and Yaser Sheikh
Carnegie Mellon University

Abstract

A core challenge in large-scale dynamic 3D reconstruction is visibility estimation---estimating which cameras observe which points at each instant in time. In this paper, we present a method to reason about the time-varying visibility of a 3D moving point captured by a large number of cameras. Our algorithm takes, as input, camera poses and image sequences, and outputs the time-varying set of the cameras in which a target point is visible. We formulate visibility estimation as a maximum a posteriori (MAP) estimate using photometric consistency, motion consistency, and geometric consistency, in conjunction with a proximal camera network prior. We demonstrate that our estimated visibility increases reconstruction performance in accuracy and density.

Dynamic 3D Reconstruction

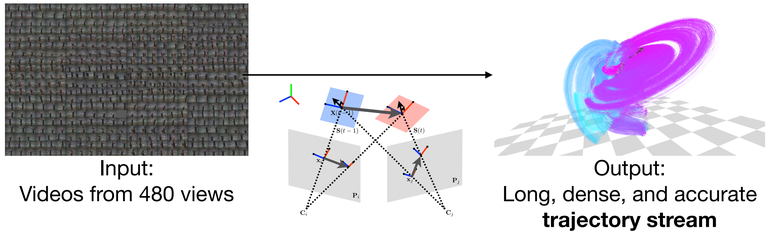

Input:
Videos from 480 views

Output:
Long, dense, and accurate
trajectory stream

Challenge

Time-varying visibility reasoning: which cameras are observing which point at each time instance?

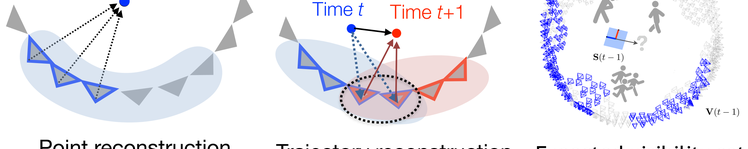

Point reconstruction

Trajectory reconstruction

Expected visibility set

Visibility Reasoning Method

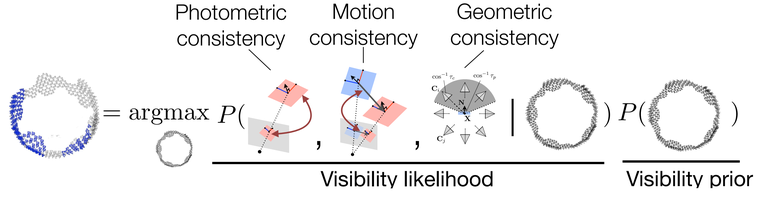

Photometric consistency

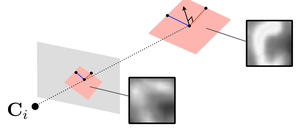

Accurate 3D patch shape and
texture of the 3D patch are required

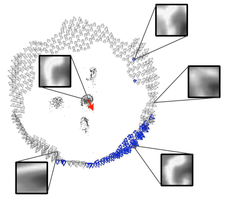

Tend to have narrow visibility set

Motion consistency

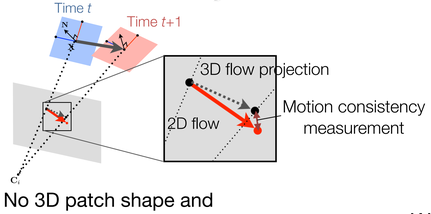

No 3D patch shape and
texture of the 3D patch are required

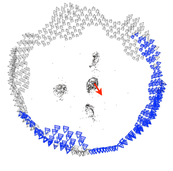

Wide and accurate visibility set

Geometric consistency

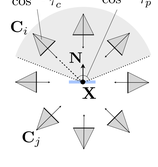

Visibility prior

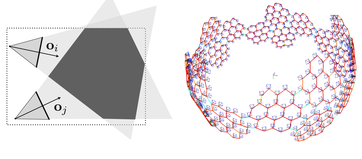

Quantitative Result

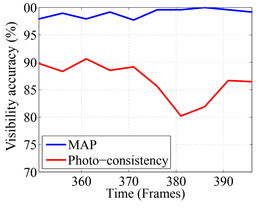

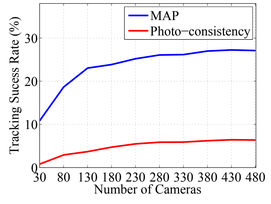

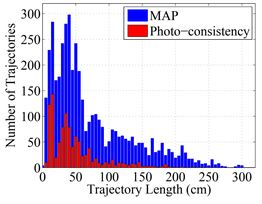

Qualitative Result