Skeletal graphs for efficient structure from motion

Supplemental material

This document contains additional visualizations of skeletal graphs and reconstructions.

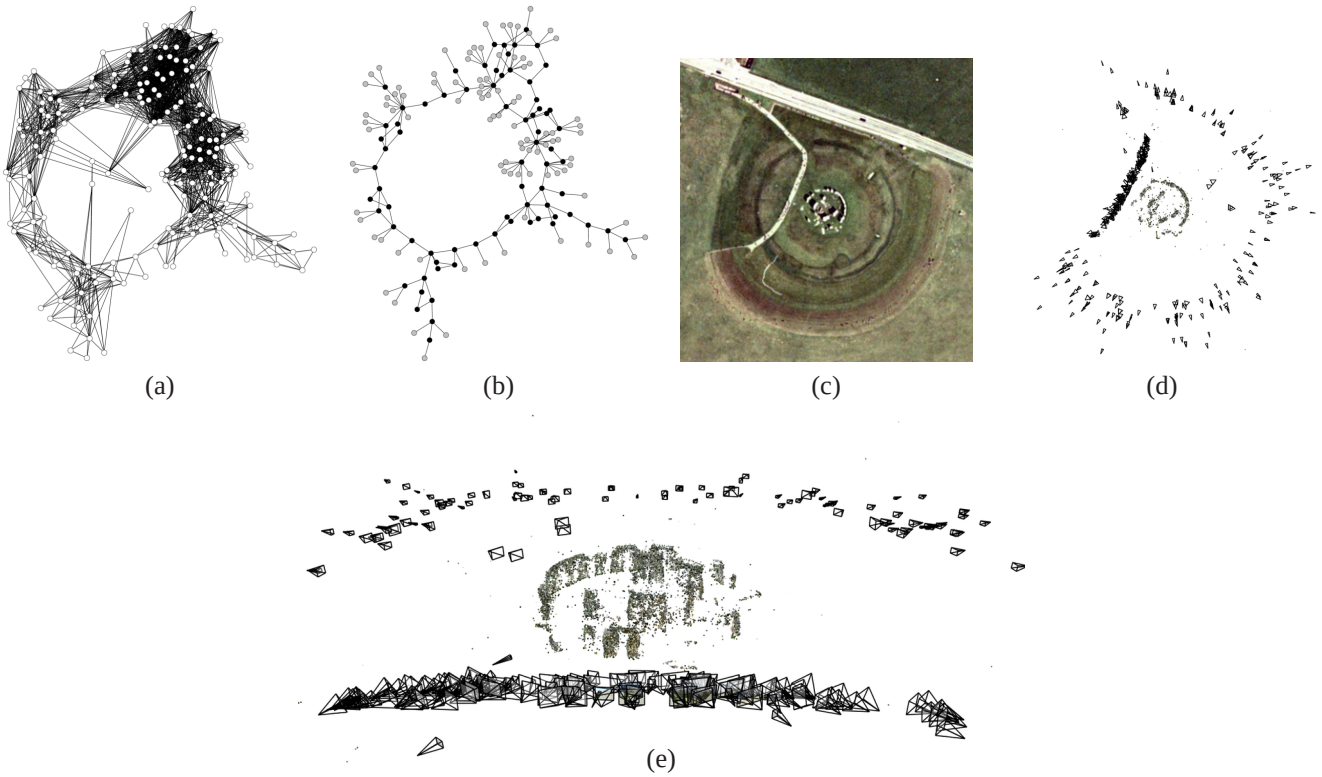

(a) (b) (c) (d)

(e)

Figure 1. *Reconstruction of Stonehenge*. (a) The full image graph for Stonehenge and (b) our skeletal graph (for $t = 16$). The black (interior) nodes of (b) comprise the skeletal set, and the gray (leaf) nodes are added in later. Note that the large loop in the graph is preserved. (The graph layouts are not based on the physical position of the cameras, but are created with the *neato* tool in the Graphviz package). (c) Aerial photo. (d-e) Two views of the Stonehenge reconstruction. The reconstructions show the recovered cameras, rendered as small frusta, in addition to the point cloud.

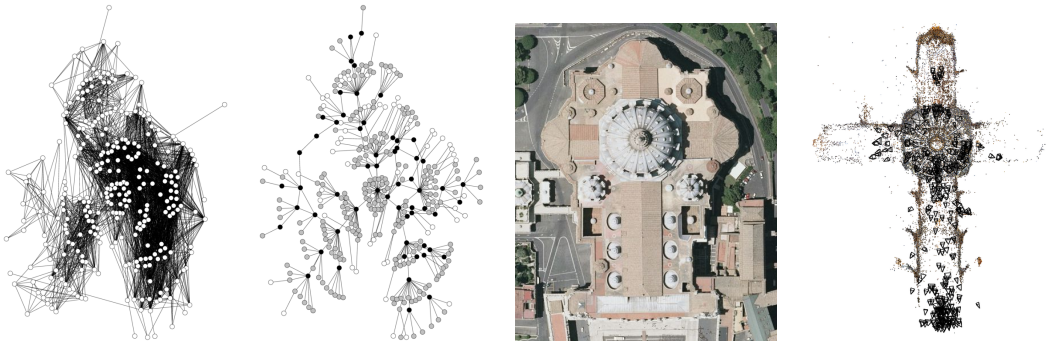

(a) (b) (c) (d)

Figure 2. *Reconstruction of the interior of St. Peter's.* (a) The full image graph for St. Peter's and (b) our skeletal graph (for $t = 16$). In this graph, white nodes represent images found to be duplicates. These nodes are removed before computing the skeletal graph. (c) Overhead photo of St. Peter's. (d) Our final reconstruction.

Figure 3. Overhead photo and view of the Pantheon reconstruction.

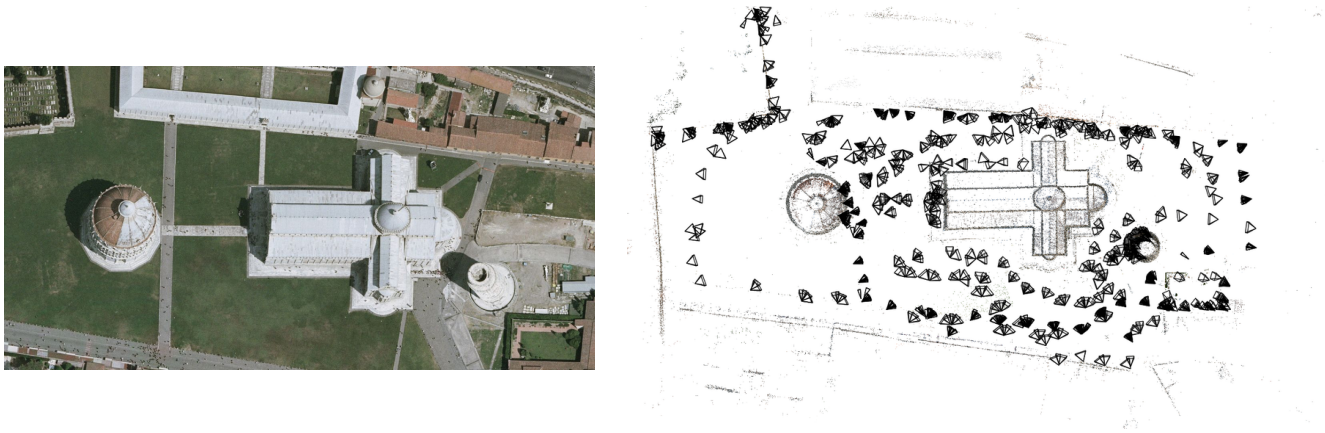

Figure 4. Overhead photo and view of the Pisa2 reconstruction.

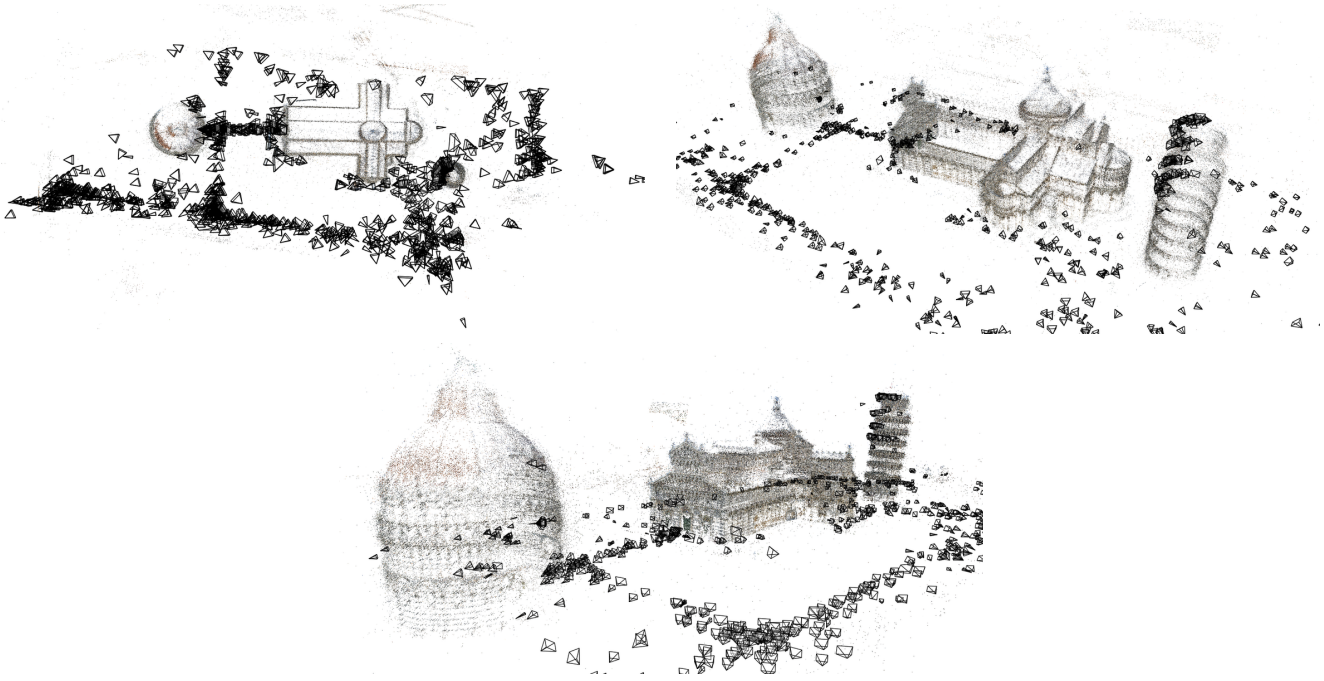

Figure 5. Several views of the Pisa1 reconstruction.

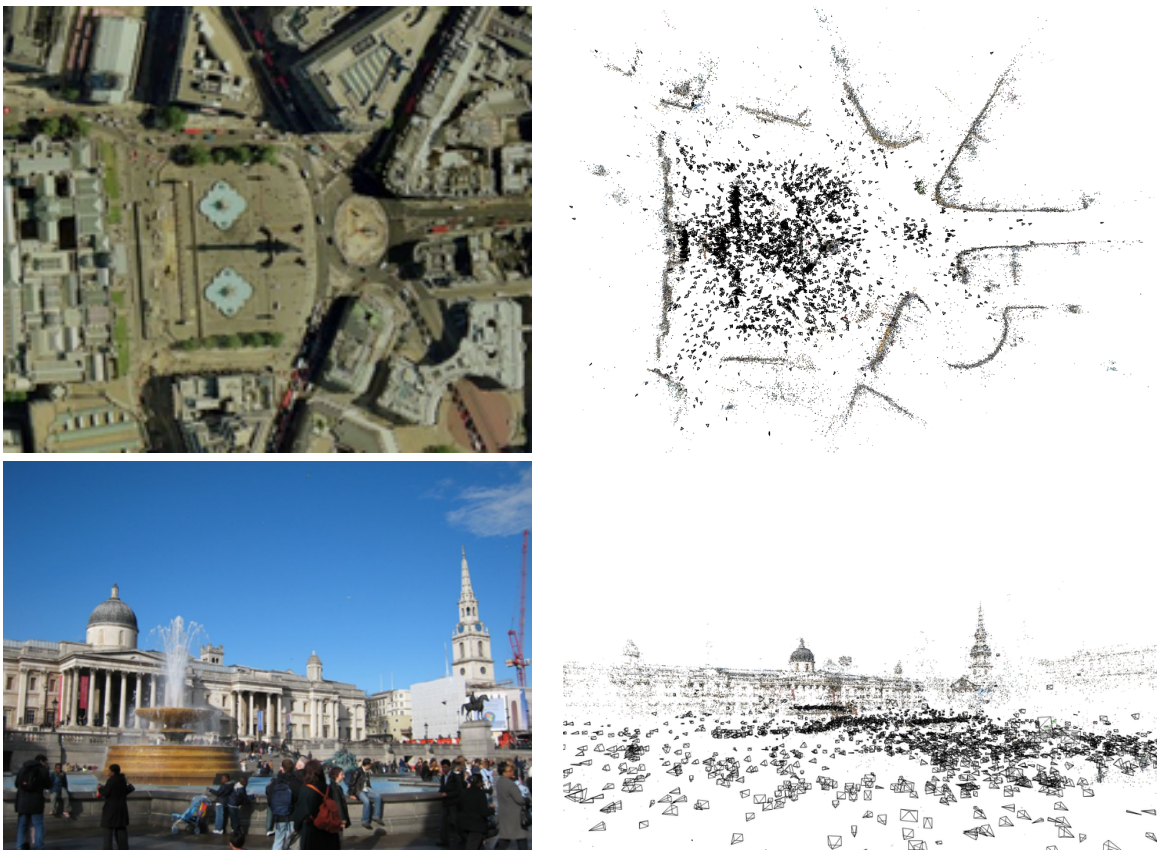

Figure 6. Views of the Trafalgar Square reconstruction, with images for comparison.

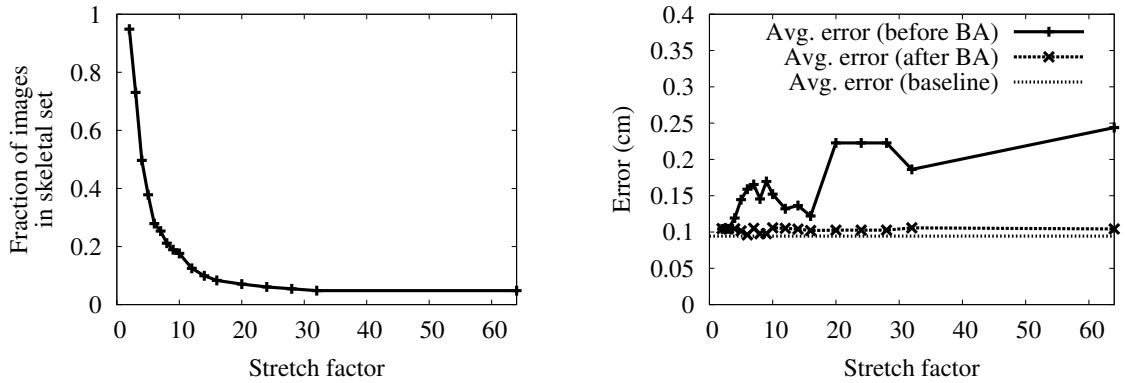

Figure 7. *Analysis of the stretch factor for the* **Temple** *data set.* This 312 image collections is taken from the multi-view stereo evaluation data of Seitz, *et al.*, and has known ground-truth camera parameters. We ran the same experiment with this dataset as with St. Peter's, described in our results. The graph on the left plots the number of images in the skeletal set for several values of the stretch factor t. The graph on the right plots the average distance between the reconstructed and ground truth camera centers for these values of t. The graph shows average error (in cm) both before and after a final bundle adjustment. Note that the error before bundle adjustment, while noisy, tends to increase with the stretch factor, but the error after bundle adjustment stays roughly constant. We also ran SfM with the full set of input images, and plot the error as a baseline. This data set is very controlled and regularly sampled compared to the Internet collections, and the behavior as t changes is somewhat different than the St. Peter's collection. For instance, even for the largest values of t we tried, the initialization provided by the skeletal reconstruction was close enough to the correct solution for the final bundle to pull it to the right place.