Finding Paths through the Worlds Photos

Tyler Nowicki

Noah Snavely, Rahul Garg, Steven M. Seitz, and Richard Szeliski. Finding paths through the world's photos. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–11, New York, NY, USA, 2008. ACM

The work presented in this paper improves photo navigation by finding paths through community photo collections (CPC). Transitioning between photos along a path is accomplished by reprojecting a nearby photo into the virtual camera (the viewer). Although the authors indicated they tried several methods for reprojection, including reprojecting from proxy geometry. They found that warping the photograph with a planar proxy produced the best results.

The quality of a reprojection is measured with three individual scores: an angular deviation score, a field of view score, and a resolution score. The individual scores are normalized to be linearly related, and multiplied together to give the reprojection score, referred to as the view score.

The angular deviation is the mean angle at each 3D point between a vector to the view position, and a vector to the camera position. Reprojections with deviations greater than 12 degrees are considered too large, so the mean angular deviation is clamped to 12 degrees. A lower angular deviation gives a higher quality reprojection.

The field-of-view score indicates how much of the display area is covered by the reprojected photo. Viewers prefer photos placed in the centre of the display area, consequently, the centre is weighted to yield higher scores (most likely with a radial gradient).

The resolution score avoids blurring by giving a low score to photos which require magnification on the display area. The score is normalized such that any value < 1 indicates the degree of magnification required for the photo.

This research constructed three different types of paths from the CPC. The orbital path allows the viewer to move in a circle around an object of interest (assumed to be in the centre of each photo). Candidate orbits are identified by finding a set of photos which are directed toward the same object of interest. Candidate orbits are thrown away if they are too short (less than 60 degrees), or if their score within a 15 degree sub interval is less than 0.01. Planar proxies do not, on their own, guarantee that the object of interest will be projected into the centre of display area. Orbit stabilization accomplishes this by rotating the proxy plane to move the object's reprojection to the centre of the display area.

Candidate panoramas are formed by identifying sets of photos taken near each other. Each panorama is scored by adding the reprojection score of each photo in the candidate set. Candidate panoramas which score below a given threshold are eliminated.

Arbitrary paths are generated when the viewer selects a photo (from a canonical set) for viewing. The path, that maximize the total view score, through a subset of the camera positions in the scene, is calculated with Dijkstra's algorithm. The sum of view scores sampled at 30 intervals between neighbouring photos in the path gives the total view score of the path.

Appearance stabilization is also applied to all reprojected photos. The reprojected photo to be display is modified with deformable registration, and colour compensation (gain and offset) to force it to resemble the previously display photo. A method is also presented for identifying night and day photos by disseminating labels to photos with matching SIFT features. This method only required a small amount of viewer input (initial photo labels), and was effective for separating different appearance states. A similar method was described in their previous work for tagging specific regions within a scene.

Finding paths is a nice generalization on previous work for finding panoramas, or other forms of photo sticking. This generalization opens up the field to many new methods of organizing and collecting photos. View scoring is the key idea that allows paths to be found within the CPC. Orbital paths were very successful, this is likely because the authors identified that viewers preferred to orbit an object of interest. Similarly, the method for finding panoramas was also successful (it is safe to assume that panoramas are also desirable). Unfortunately, the method for identifying arbitrary paths was not inspired by how viewers prefer to navigate a scene or set of photos. View scoring should not have been replied upon to produce paths that the viewer would find interesting.