Three-Dimensional Kaleidoscopic Imaging

Ilya Reshetouski, Alkhazur Manakov, Hans-Peter Seidel, Ivo Ihrke

Saarland University / MPI Informatik
campus E1.7, 66 123 Saarbrücken, Germany
{*iresheto*|*amanakov*|*ihrke*}*@mmci.uni-saarland.de*

Abstract: Planar mirror systems are capable of generating many virtual views, yet their practical use for multi-view imaging has been hindered by limiting configurations that enable view decomposition. In this work we lift those restrictions.

OCIS codes: 110.1758, 150.6910.

1. Introduction

Since its invention by David Brewster in 1815 the Kaleidoscope has fascinated our minds. Its ability to generate hundreds of intricately interwoven views of the same object generates beautifully patterned images. For imaging purposes, kaleidoscopic systems have so far been used for reflectance measurements [1, 2] owing to their ability to generate a large number of views of the target, almost covering the full hemisphere surrounding the measured surface patch. However, state of the art systems are restricted to imaging flat samples [1,2], or virtual views that do not overlap on the image plane [3–5].

The limitation of only being able to image planar objects is due to the fact that it is a-priory impossible to determine the virtual view of any one pixel if the object geometry is unknown. A major challenge is potential self-occlusion of the object that is being imaged. In this report, we propose a general solution to this problem.

2. Kaleidoscopic Imaging Theory

2.1. Space Partitioning

The goal of this subsection is to introduce the tools necessary to understand image formation in complex systems of planar mirrors. For convenience of illustration we develop the concepts in two dimensions, the generalization to the three-dimensional case is straight-forward.

Consider a pinhole camera with projective center C in the mirrors triangle (base chamber) of a kaleidoscopic imaging system. Our goal is to describe the resulting reflected ray geometry, see Fig. 1 (a), in an intuitive way. The ray l is

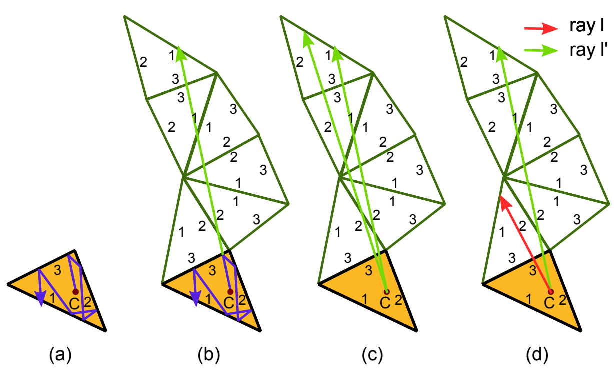

Fig. 1. A ray in a kaleidoscopic system is reflected off the planar mirrors (a). "Unfolding" of the ray can be performed by mirroring the base chamber instead of the ray (b). Two neighboring rays typically share a common unfolding scheme (c). Upon intersection with different mirror planes this coherence breaks down (d).

reflected off the mirrors and traverses them in a particular order $L_N = (3,2,1,2,1,3,1,\ldots)$ for a particular number N of reflections. Here, L_N is an ordered N-tuple describing the mirror sequence that the ray intersects. It is indicative of the light path taken by the ray inside the system. We can "unfold" this light path by mirroring the base chamber instead of the ray, creating a sequence of mirror chambers along the straightened ray, see Fig. 1 (b). Upon transforming the mirror chambers, along with the ray segment contained therein, back to the base chamber, we re-obtain the "folded" light path of Fig. 1 (a).

In general, two neighboring light paths l and l', Fig. 1 (c), traverse the system in a similar manner, i.e. $L_N = L'_N$ for some reflection count N. This is the reason for obtaining recognizable virtual views in systems of planar mirrors. As long as this condition persists, the space partitioning generated by one of the rays, e.g. Fig. 1 (c), is a valid explanation for both l and l'. This argument breaks down when the two rays hit different mirrors at some reflection count $\hat{N} > N$, see Fig. 1 (d), and $L_{\hat{N}} \neq L'_{\hat{N}}$. This causes discontinuities in the space partitioning, but the complete space remains covered without overlap.

2.2. System Containing Objects

Ultimately, our goal is to label each pixel of our projective camera image with the mirror chamber where its corresponding ray intersects the object. This information corresponds to determining the virtual view point for each pixel of the kaleidoscopic view. An illustration of this is shown in the Fig. 2. Differently colored cones encode different viewpoints of the object. Objects $1-4$ are fully visible, 5 and 6 are partially visible while 7 is fully occluded. Even if 7 was visible, only part of it could be observed due to the space partitioning introduced by the mirrors. To achieve

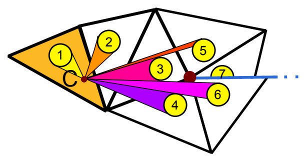

Fig. 2. Camera observing the object in the kaleidoscope.

the labeling, it is necessary to infer some geometric structure of the object under consideration. As can be seen from the figure, once the object geometry is known, labeling the pixels is trivial. The object only has to be projected to the camera center C from all its visible mirrored positions, as determined by the space-partitioning of the system, with occlusion taken into account.

Our method to determine an approximate geometry of the object is based on considering rays that do *not intersect* any real or mirrored version of the object.

2.3. Reconstructing the Visual Hull of the Object

It is possible to design a setup where gaps between an object and its mirror images are observable. These gaps provide the means to perform a visual hull [6] reconstruction of the object.

Consider a ray that does not intersect the object nor any of its virtual counterparts. Folding back this ray into the base chamber, we obtain a reflected light path that is guaranteed to be free of intersections. Performing this operation on the set of all rays that do not intersect the object anywhere in mirror space, we obtain a space carving scheme to determine the visual hull of the object: It is computed by successively removing free space from an initial volume that is marked as containing the object.

2.4. Labeling the Image

The visual hull computed this way can effectively be used as a geometric proxy for the object. By transforming the visual hull into the mirror chambers and intersecting the straightened camera rays in unfolded mirror space with the set of mirrored visual hulls, we can label the rays w.r.t. the mirror chamber where the ray first intersects the visual hull. This way, a virtual view of the object is determined for each pixel.

3. Practical implementation

3.1. Kaleidoscope design choice

We chose the frustum of a triangular pyramid as our base chamber, the narrow end pointed downward. In this type of setup the views are arranged spherically around the object, resulting in a high view point variation.

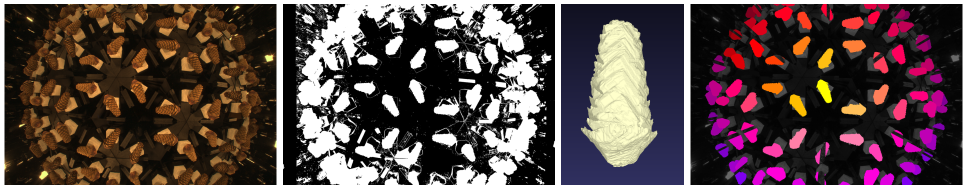

Fig. 3. Results (left to right) input image, silhouette image, computed visual hull, corresponding labeling.

# refl.	subsample	VH disc.	# virt. views	# labeled pixels
7	9	$256 \times 256 \times 236$	128	18.38%
8	9	$256 \times 256 \times 236$	166	19.60%
9	9	$256 \times 256 \times 236$	212	19.91%

Table 1. Statistics for the Cone data set. From left to right: number of reflection levels used to compute result, number of sub-samples per pixel, discretization of the visual hull, number of virtual views that have been used for computation and the number of labeled pixels.

3.2. Experimental Results

We recorded our images at a resolution of 3866×2574 pixels. The data sets as well as the computed labeling images are shown in Fig. 3. Table 1 summarizes some statistics. All results have been computed only using the silhouette image shown in Fig. 3 and the calibration information. The results in the figure were computed using the information from 8 levels of reflection, equalling 166 views. Further results can be found on our web page: http://giana.mmci.uni-saarland.de/projects/kaleidoscopic_imaging/.

4. Summary and Future Work

We have introduced a general framework for dealing with systems of planar mirrors imaged by a projective camera. We have shown that generalized kaleidoscopic imaging systems can be used to obtain dense hemispherical multi-view data that is calibrated both geometrically and photometrically. The output of our techniques is thus directly usable in standard multi-view reconstruction algorithms. Due to the wide range of views achievable with these systems it is possible to image every surface point from a large number of directions. It might thus become possible to perform simultaneous geometry and reflectance estimation on dynamic objects.

Future work includes the incorporation of the inherently multi-resolution data generated by our system into multi-view reconstruction algorithms as well as investigating techniques to differentiate between the limited number of views that are potentially responsible for unreliable pixels.

References

1. J. Y. Han and K. Perlin, "Measuring Bidirectional Texture Reflectance with a Kaleidoscope," in "Proc. SIGGRAPH," (ACM, New York, NY, USA, 2003), pp. 741–748.
2. S. Bangay and J. D. Radloff, "Kaleidoscope Configurations for Reflectance Measurement," in "Proc. AFRIGRAPH," (2004), pp. 161–170.
3. K. Forbes, F. Nicolls, G. D. Jager, and A. Voigt, "Shape-from-Silhouette with two Mirrors and an Uncalibrated Camera," in "Proc. ECCV," (2006), pp. 165–178.
4. P.-H. Huang and S.-H. Lai, "Contour-Based Structure from Reflection," in "Proc. CVPR," (2006), pp. 165–178.
5. D. Lanman, D. Crispell, and G. Taubin, "Surround Structured Lighting for Full Object Scanning," in "Proc. 3DIM," (2007), pp. 107–116.
6. A. Laurentini, "The Visual Hull Concept for Silhouette-Based Image Understanding," IEEE Trans. PAMI pp. 150–162 (1994).