

# Supplementary Material

Discovering the Structure of a Planar Mirror System  
from Multiple Observations of a Single Point

## 1 $2\frac{1}{2}$ D synthetic example.

This section discusses how our two-dimensional theoretical results can be applied to the  $2\frac{1}{2}$ D case, a situation often occurring in practice. We discuss the synthesis aspect of a simulation example in some detail to illustrate the principal reasoning when moving from  $2\frac{1}{2}$ D to  $2D$  (in order to apply our algorithms).

In our synthetic experiment we simulated the model of a  $2\frac{1}{2}$ D kaleidoscope with 6 planar mirrors (Fig. 1), a one-sided planar checkerboard (Fig. 2), and a camera with a 50mm lens.

The camera was positioned near the entrance hole of the kaleidoscope and a checkerboard was placed inside. The image, observed by the camera is shown in Fig. 3. The viewing direction of the camera and the position of the checkerboard is selected such, that the camera cannot observe the checkerboard directly. The minimum reflection level among all observed checkerboards is two.

For data generation, we exploit the known pose of the camera, the checkerboard, and the mirrors. With this information we can construct ground truth 3D coordinates of some of the central points of the virtual checkerboards observed by the camera, i.e. we simulate a perfect measurement setup.

In practice, we use the known 3D position of the real checkerboard center and generate all possible virtual copies up to a certain number of reflections, Fig. 4. Note, that all these 3D points are arranged in parallel layers (see Fig. 5). For identifying the visible points in our simulated view, we back-project the image points of the virtual checkerboard centers marked by a yellow disk in Fig. 6 and identify the simulated 3D point closest to the back-projected ray. The results of this operation are shown in Figs. 7 and 8. Using this set of selected 3D points, and exploiting the knowledge that the geometry of our kaleidoscope results in points that are arranged in equidistant planar layers parallel to the top (or bottom) mirror plane, we can transfer to a 2D case by projecting all the selected 3D points onto the top mirror plane. The selected 3D points are now transferred to 2D and we can work with the algorithms developed in the paper. The 2D versions of the selected points and the camera center for the synthetic example are shown in Fig. 9.




Figure 1: A model of the synthetic kaleidoscope. View from outside. It has an entrance hole for the light source and the camera. We are aiming at recovering the convex hull of the two-dimensional ground plan of the structure, i.e. we want to reconstruct the side-mirror planes as well as the camera pose.

We perform the reconstruction of the mirror geometry in 3 steps: construction of the full doublet graph, Fig. 10, filtering of the doublet graph Fig. 11, and randomized forward search Fig. 12. The results perfectly match the given initial model. In Table 1, we provide statistics for the results discussed in this section. The large number of trials can be explained by the very small field-of-view of  $\approx 40^\circ$  that the simulated  $50mm$  system imposes, see also Sect. 3 for more insight on this issue.

Note that as our checkerboard is one-sided, we are only observing odd levels of reflections from the top and bottom mirrors. Thus the distance between the layers is a multiple of the double distance between the top and the bottom mirrors. The position of the observed odd reflections of the checkerboard is invariant with respect to a rigid shift of the pair of top and bottom mirrors along their normal. Therefore, information on the placement of the camera inside the kaleidoscope can only be recovered up to an unknown shift of the camera center along the top mirror normal. If this information is required, we fix it by manually picking one of the recovered lines of discontinuity to agree

#iter	1910	2202	2781	3714	4549	5152	6807	7136	7596	9681
time(s)	24.3	29	38.59	56.13	71.4	83.8	124.8	132.8	147.7	203.5

Table 1: Statistics for 10 different runs for recovering the  $2\frac{1}{2}$ D synthetic example. We show the number of cycle discovery restarts and the time required to compute the solution.

with its recorded counterpart.

## 2 $2\frac{1}{2}$ D real mirror example.

This section gives some more details about the real world experiment discussed in the paper.

In our real experiment we set up a kaleidoscope with 6 planar mirrors resembling the configuration of the synthetic experiment. For capturing the images we used a Canon 5D mark II equipped with a  $35mm$  lens. Fig. 13 shows our input image, taken in the kaleidoscope. The camera’s intrinsics were calibrated with Bouguet’s calibration toolbox for MATLAB. We extract the center points of the virtual checkerboards automatically by using fiducial markers (example, Fig. 14). We triangulated the selected points by estimating homographies using neighboring marker positions, Fig. 15. The fiducial markers were also used to reject incorrect neighbors for the homography estimation.

To estimate the top (and bottom) mirror plane orientations we fit a system of parallel equidistant planes to the selected 3D points, Fig. 16. After projecting all the points onto a common plane as in the synthetic example, we again reduce the problem to a 2D scenario. Figs. 17, 18, and 19 are illustrating the process of solving the 2D reconstruction problem. After reconstructing the ground plan of the side mirrors and computing the distance between the top and the bottom mirrors as half the distance between nearest layers of the selected 3D points, we have recovered an estimation of the complete 3D model of the kaleidoscope, the 3D position of the center of the real checkerboard, and the pose of the camera up to one degree of freedom (the camera and the real checkerboard could be at any possible position along the vertical axis (perpendicular to the top mirror plane), where the checkerboard position is uniquely determined by the camera position). In Fig. 20 we show a result where we superimpose the camera image and the reconstructed mirrors.

## 3 Additional Statistics

In this section we provide additional statistics of our reconstruction scheme. The results should be interpreted as those in Fig. 5 in the paper. We investigated the behavior of the randomized search algorithm with respect to the allowed maximum number of cycle discovery restarts (iterations). Figs. 21 and 22 show the results. Again, the exhaustive search algorithm on the clean doublet graph

serves as a baseline. If this algorithm cannot find a solution, any triplet-based algorithm will fail. The randomized forward search algorithm, however, has to discover the solution in the filtered doublet graph that contains false connections. The figures show the number of successful reconstructions out of 1000 randomly generated cases that can be achieved with a number of less than or equal to  $\{2, 4, 8, 16, 32, 64, 128, 256\}$  cycle discovery restarts. In the full surround view case, the curves converge to a limit curve. Increasing the number of iterations does not increase our chances of success. In the limited field-of-view case, the situation improves linearly with an exponential increase in the iterations. The results suggest that improved graph exploration techniques may improve recoverability towards the base line curve. The plots motivate our choice of 50 iterations for the synthetic experiments conducted in the paper.




Figure 2: Synthetic configuration: top and bottom sides of the simulated checkerboard.




Figure 3: Synthetic configuration: a simulated view inside the kaleidoscope (50mm focal length).




Figure 4: Synthetic configuration: we show all possible virtual checkerboard center positions up to 7 levels of reflection (neglecting occlusion).




Figure 5: Synthetic configuration: this point of view shows that all the virtual checkerboard center positions are arranged in layers parallel to the top and bottom mirror.




Figure 6: Synthetic configuration: yellow disks marking eleven selected camera image points showing the centers of virtual checkerboards.






Figure 7: Synthetic configuration: 3D points corresponding to the eleven selected centers of virtual checkerboards plotted with all other possible virtual checkerboard center positions (no occlusion).




Figure 8: Synthetic configuration: From this point of view it is clear, that the real checkerboard is not among the selected points and that all selected points are to be found in two different layers.




Figure 9: Synthetic configuration: selected 3D points and the camera center projected onto the top mirror plane.




Figure 10: Synthetic configuration: the full doublet graph indicating all potential doublets (green segments) for the eleven selected 2D points (blue dots). C - position of the camera center.




Figure 11: Synthetic configuration: the filtered doublet graph, indicating all doublets left after filtering (green and red segments) for our eleven selected 2D points (blue dots). Red segments are indicating the valid doublets among them (found after randomized forward search).  $C$  - position of the camera center.




Figure 12: Synthetic configuration: final result with the reconstructed base chamber (yellow polygon), camera center  $C$  (red dot) and a corresponding unfolding up to 7 levels of reflections. We also show the eleven selected 2D points (blue dots) and the reconstructed 2D position of the checkerboard center  $S$  (gray dot). As can be seen, the field-of-view of the simulated system is rather low ( $\approx 40^\circ$ ).



Figure 13: Real world example: view inside the mirror system.



Figure 14: Real world example: selected image points (red dots), corresponding to centers of virtual checkerboards.




Figure 15: Real world example: 3D positions of the selected image points.






Figure 16: Real world example: selected 3D points are grouped into four pairwise-parallel planar layers, with distances between them as multiples of the double distance between the top and the bottom mirrors.





C

Figure 17: Real world example: the full doublet graph indicating all potential doublets (green segments) for the fifteen selected 2D points (blue dots). C - position of the camera center.



C

Figure 18: Real world example: the filtered doublet graph, indicating all doublets left after filtering (green and red segments) for our fifteen selected 2D points (blue dots). Red segments are indicating the valid doublets among them (found after randomized forward search). C - position of the camera center.

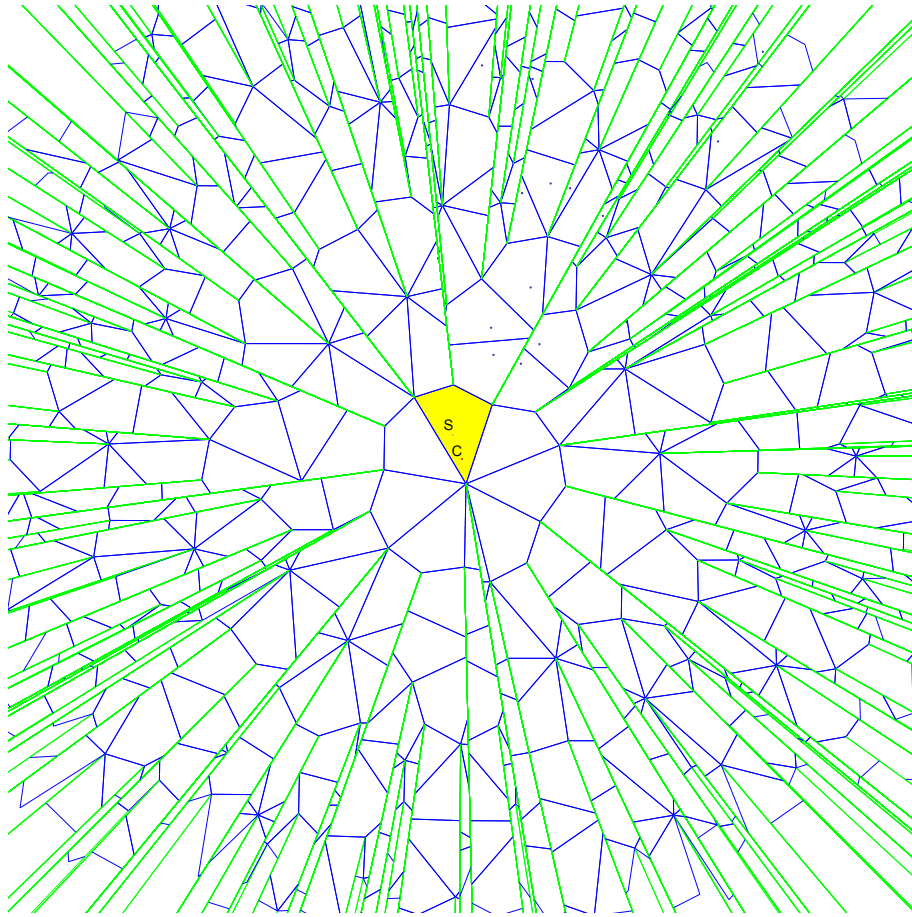


Figure 19: Real world example: final result with the reconstructed base chamber (yellow polygon), camera center  $C$  (red dot) and a corresponding unfolding up to 11 levels of reflections. We also show the fifteen selected 2D points (blue dots) and the reconstructed 2D position of the checkerboard center  $S$  (gray dot).




Figure 20: Real world example: mirrors are super-imposed on the image as attenuating layers.

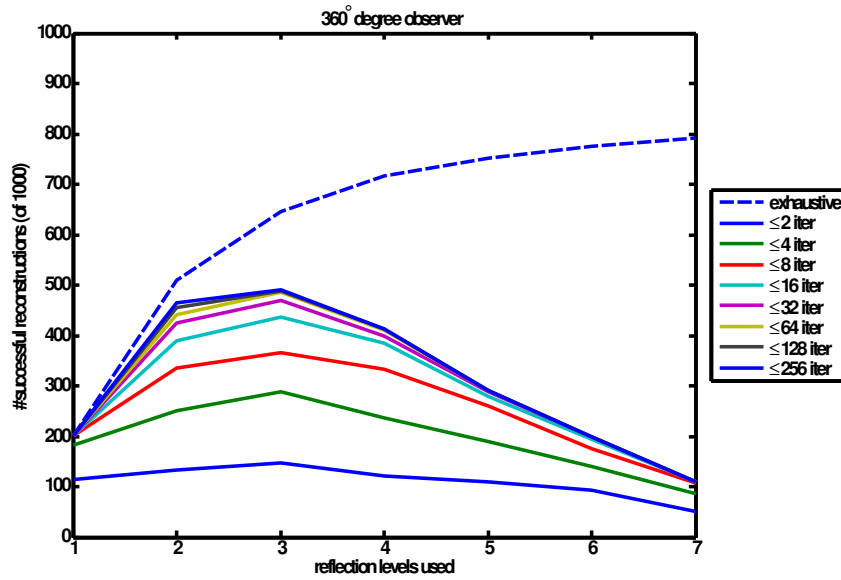


Figure 21: Additional simulations results: this plot indicates the number of iterations until successful recovery of the geometry using the randomized search algorithm for a  $360^\circ$  observer. The dashed line shows the base line provided by the exhaustive search algorithm running on the clean doublet graph. This is the maximum number of achievable reconstructions in 1000 exemplars (in the remaining cases there is insufficient information for reconstruction even in the perfect case). The randomized algorithm runs on the imperfectly filtered doublet graph. The number of cycle discovery trials follow a logarithmic distribution. The discovery rate converges towards a limit curve.

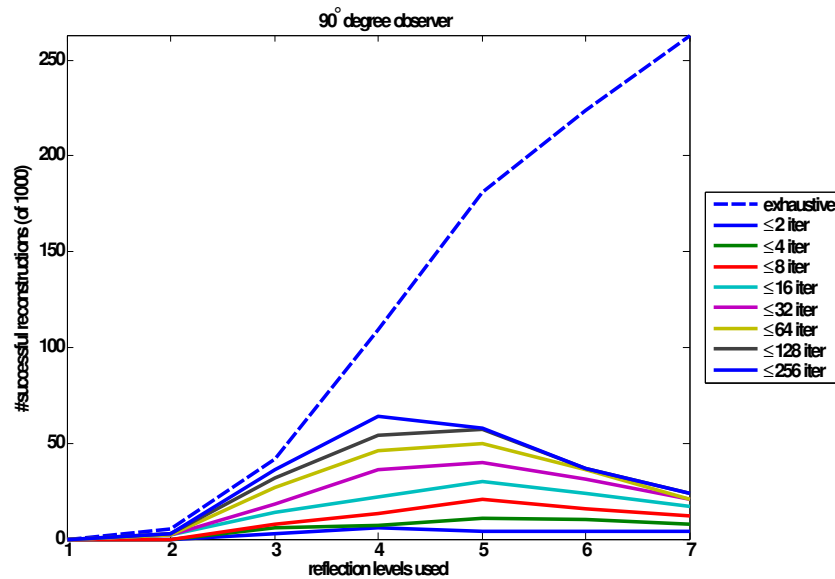


Figure 22: Additional simulations results: this plot indicates the number of iterations until successful recovery of the geometry using the randomized search algorithm for a  $90^\circ$  observer. The dashed line shows the base line provided by the exhaustive search algorithm running on the clean doublet graph. This is the maximum number of achievable reconstructions in 1000 exemplars (in the remaining cases there is insufficient information for reconstruction even in the perfect case). The randomized algorithm runs on the imperfectly filtered doublet graph. The number of cycle discovery trials follow a logarithmic distribution. The discovery rate converges towards a limit curve, but more slowly than in the full surround case. Please note that the scale of the plot is magnified to show details of the simulation curves (the maximum number of recoverable geometries is  $\approx 250/1000$  in this case.)