A Flexible Projector-Camera System for Multi-Planar Displays

Mark Ashdown
mark@ashdown.name
University of Tokyo Institute of Industrial Science

Matthew Flagg
mflagg@cc.gatech.edu
College of Computing, Georgia Institute of Technology

Rahul Sukthankar
rahuls@cs.cmu.edu
Intel Research Pittsburgh & The Robotics Institute, Carnegie Mellon

James M. Rehg
rehg@cc.gatech.edu
College of Computing, Georgia Institute of Technology

Poster for CVPR 2004

Introduction

We present a novel multi-planar system based on an uncalibrated projector-camera pair. In contrast to previous room-level projector-camera systems, our method is based on a flexible calibration procedure that requires minimal information about the geometry of the multi-surface scenario. It creates *ad hoc* visualization and display capabilities on surfaces in a home or office environment.

The key to our calibration approach is an efficient technique for simultaneously localizing multiple planes and a robust planar metric rectification method that can tolerate a restricted camera field-of-view and requires no special calibration objects.

Calibration consists of three stages. The first is the identification of planar surfaces and recovery of homographies from projector to camera through each surface. The second performs metric rectification and alignment to obtain homographies from the camera to each surface. The third stage adjusts the homographies from projector to each surface to reflect the constraints imposed by the intersection of pairs of surfaces.

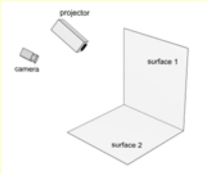

The uncalibrated projector and camera are casually positioned in front of the display surfaces

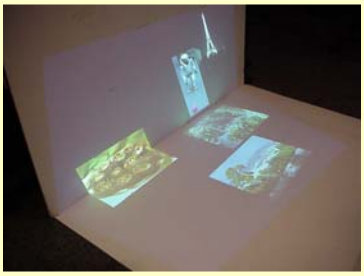


A photograph is dragged from one plane to another in this photo-browser application. The two surfaces are spanned by a single projector.

A user at a two-surface projected interface. The projector and camera can be mounted at the side.

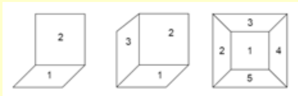
Three multi-planar configurations: a desk placed against a wall, a table pushed into the corner of a room, and the inside of a box. There will be a projector-to-surface homography for each surface.


The projector-to-surface homographies are not independent. The graphs above show the constraints between them, where the nodes are homographies and the edges are the constraints.

Identifying Planar Surfaces

The goal of the first stage of calibration is to automatically detect, segment, and calibrate a piece-wise planar scene into a set of connected surfaces. We use an uncalibrated variant of structured light: the projector displays a series of horizontal and vertical lines that are observed by the camera. A line that crosses multiple surfaces appears as several line sections.We fit a line to each section, intersect the line to find *kinks* and, fit a line to the kinks to determine the boundary between surfaces. Using this boundary we can process the sets of line sections on the two surfaces separately.

For each surface i we calculate a planar homography from projector to camera through the surface. Using lines of the form

$$\mathbf{l} = (a, b, c)^\top$$

the constraints

$$\mathbf{l}'_i \times \mathtt{H}^{-\top}\mathbf{l}_i = \mathbf{0}$$

form a closed-form least squares solution similar to that from point correspondences. We achieve robustness to outliers by using the closed-form solution on a series of random minimal subsets of the line correspondences, picking the subset that best fits the data using the median of an algebraic distance measure, then discarding the correspondences that do not fit the selected homography.


Lines are displayed by the projector. *Kinks* in the projected lines indicate the location of the boundary between the surfaces

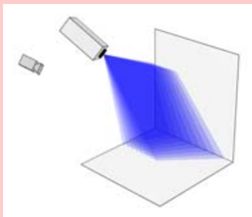
The projector displays multiple lines that allow the boundary between planes to be detected.

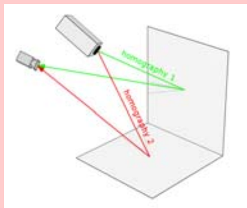

We calculate a projector-to-camera homography through each plane using line correspondences between projector and camera

Metric Rectification

The next step is to recover the homographies from the camera to each surface. We decompose the homography into two parts: a metric-rectification that maps the image of each surface to an arbitrary Euclidean frame, and a similarity transform that aligns the frame to the physical surface. To obtain the former the camera observes a set of arbitrarily placed right angles on each surface. These are obtained by imaging everyday objects such as postcards.

The conic dual to the circular points on a Euclidian surface is

$$\mathtt{C}^*_\infty = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

For perpendicular lines **l** and **m**

$$\mathbf{l}^\top \mathtt{C}^*_\infty \mathbf{m} = 0$$

When a surface is imaged by the camera it is transformed by a planar homography and the conic dual becomes

$$\mathtt{C}^{*\prime}_\infty = \mathtt{H}\mathtt{C}^*_\infty\mathtt{H}^\top$$

but for the images **l**´ and **m**´ of the lines we still have

$$\mathbf{l}'^\top \mathtt{C}^{*\prime}_\infty \mathbf{m}' = 0$$

thus images of lines that are perpendicular on the surface place linear constraints on the transformed conic dual matrix. We use a closed-form least-squares method to calculate the matrix.

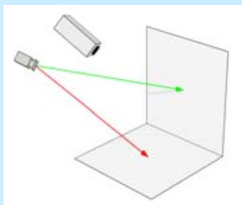
We wish to recover the homography from the camera to each surface

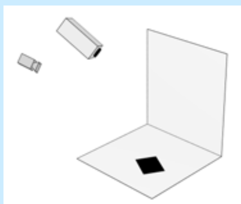
An everyday object that contains right-angles, like a post card, is placed on the surface

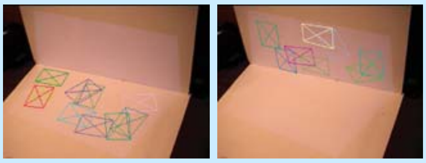

Orthogonal line pairs extracted from randomly placed postcards are used to achieve metric rectification.

To obtain an estimate for the camera-to-surface homography we take the singular value decomposition of the (symmetric) transformed conic dual matrix to get the form

$$C_\infty^{*\prime} = UDU^\top$$

Assuming the last singular value to be small, we set it to zero and express the diagonal matrix as

$$D = BC_\infty^* B^\top$$

where

$$B = B^\top = \begin{bmatrix} \pm\sqrt{D_{11}} & 0 & 0 \\ 0 & \pm\sqrt{D_{22}} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

thus

$$U(BC_\infty^* B^\top)U^\top = HC_\infty^* H^\top$$

$$H = UB$$

The choice of signs in **B** creates four possibilities for **H** but only two need to be considered. We determine the correct one by testing the possibilities with the real data.

The homography **H** will now allow metric rectification of images of the surface but to complete the homography from camera to surface we require a similarity transform to align the projected graphics with the physical surface. We obtain this by constraining the origin of each surface's co-ordinate frame to be on the boundary and the x-axis to lie along that boundary, and selecting a scale that is consistent between surfaces.

The SVD estimation for metric rectification can suffer from normalization and robustness problems in practice. We have found that centering and sphering the line data dramatically improves the conditioning of the problem and should be considered essential.

As with the line homography case, to achieve robustness to outliers we compute the metric rectification with minimal subsets of the constraints and test those solutions on the remaining line correspondences. In this case, 5 correspondences are required for each subset to obtain the 3 by 3 symmetric conic dual matrix up to an insignificant scale factor.

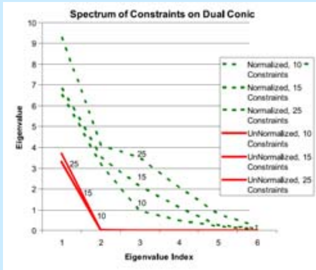

Normalization of the data from the lines detected on the surfaces improves the accuracy of the metric rectification. This graph shows the spectrum of the constraints with and without normalization.

Constraints Between Homographies

We calculate the projector-to-surface homographies independently so there is no guarantee that they will be consistent along the boundaries between surfaces. We therefore provide an iterative refinement algorithm to enforce the inter-surface constraints.

For two surfaces we start with two surface-to-projector homographies

$${}^1_S H_P \quad \text{and} \quad {}^2_S H_P$$

For points on the boundary between surfaces we would like the corresponding projector points to be equal, that is

$${}^1_S H_P \mathbf{x} = {}^2_S H_P \mathbf{x}$$

But this will typically not be true due to image processing error, so we apply the constraint by refining the homographies with an iterative algorithm.

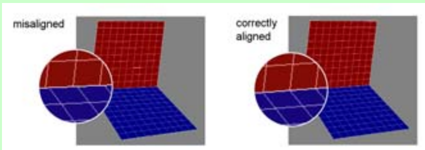

Refinement of the projector-to-surface homographies is necessary to ensure that the co-ordinates on the surfaces match up along the boundary.

The diagram on the right shows the steps of the algorithm. We start by generating five points on each surface (a) and generate the initial correspondences in the projector using the current values of the two homographies (b). The locations of the points on the surfaces are fixed throughout the algorithm while their corresponding points in the projector are adjusted. At least three points are needed on the shared boundary to ensure that length ratios agree along the boundary.

During the iteration, each point on the boundary is transformed to two points in the projector using the two homographies. A point in the projector is then created at the midpoint of those two (c), and two homographies are recomputed using the new point correspondences thus generated.

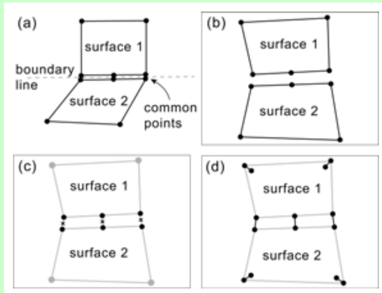

The steps of the homography refinement algorithm. Points on the surfaces are kept constant while their correspondences in the projector are varied.

The error for the homographies is measured as the the sum of the separations of between desired and transformed locations in the projector frame (d), and the algorithm terminates when this error falls below a threshold.

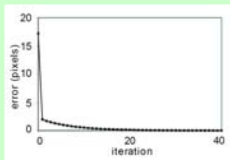

The misalignment between homographies is quickly reduced.

Conclusions & Future Work

We have presented robust calibration algorithms for aligning a camera-projector system to multiple planar surfaces. The algorithm for calculating homographies from line correspondences works even when a significant fraction of the data consists of outliers, which enables us to deploy the uncalibrated structured light system in difficult imaging situations. The robust metric rectification algorithm employs everyday objects such as postcards, to determine the camera-to-surface homography for each of the multi-planar surfaces. The iterative homography refinement technique reduces inconsistencies along the boundaries of multi-surface displays.

We have implemented an interactive application for a multi-surface display that enables users to manipulate images on each surface and move them between surfaces. Multi-surface displays could also be used to present co-ordinated views of 2D data such as the plan and elevation views of an architectural model, or for 3D visualization by clipping or projecting voxel data onto the surfaces.

Various new visualizations could exploit a multi-planar display.

東京大学 THE UNIVERSITY OF TOKYO
GVU center georgia tech
THE ROBOTICS INSTITUTE
Intel Research laboratory @ pittsburgh
http://www.cc.gatech.edu/grads/f/Matt.Flagg