

# Projecting Rectified Images in an Augmented Environment

Stanislaw Borkowski, Olivier Riff and James L. Crowley  
email: {Stan.Borkowski, Olivier.Riff, James.Crowley}@inrialpes.fr

Projet PRIMA, Lab. GRAVIR-IMAG  
INRIA Rhône-Alpes, France

## Abstract

*Steerable video projection systems make it possible to use any convenient surface as a display. This paper describes a steerable projection system in which a projector and a camera are mounted in an overhead motorized assembly with two degrees of freedom. Associating a remote camera to the video projector makes it possible to discover planar surfaces in the environment that can be used as display surfaces.*

*One possible application of this system is a "portable display surface" composed of a hand-held white cardboard surface with a black boundary. An estimation technique is described in which the display boundaries are tracked using a robust tracker to provide the necessary information to compute an inverse homography for pre-warping the image. Continually updating the inverse homography enables the system to maintain a rectified image on the display surface despite 3D translations and rotations.*

**Keywords:** Ubiquitous Computing, Augmented environment, Motorized video-projector, Projector-scene calibration, Portable display surface.

## 1. Introduction

Associating a video projector with a video camera offers an inexpensive means to transform any surface in an environment into an interactive display surface. Rigidly mounting the video camera on the projector provides a number of useful advantages. Observing the position of points in the scene makes it possible to estimate the inverse transformation required to pre-warp an image before projection. Observing the interaction of hands or limbs with projected information provides a simple and intuitive mode for man-machine interaction.

An early system for augmenting an environment by combining a projected information and computer vision was the "digital desk" described by Wellner and MacKay in [10]. Over the last few years, several augmented environments had been designed [5, 8] using single or multiple video-

projectors to augment ordinary objects and surfaces. Until recently, such systems have used statically mounted projectors, thus restricting display to limited parts of the environment.

A device for redirecting the projector beam by a steerable mirror, overcoming the spatial constraints of fixed-projector systems, has been presented in [4]. In this system, the ability to project images in the environment is restricted to a cone resulting from the mirror-projector configuration. In this paper we describe a Steerable Video Projector platform (SVP), conceptually similar to [2], that is able to project images in any direction within a hemisphere around the device. While somewhat bulky, our device anticipates the current trend of projectors to become portable devices, similar in shape to hand-held torch lamps [1].

Steerable video projectors offer an inexpensive means to transform any convenient surface into an information display. However, unless a surface is perpendicular to the projector beam, the resulting image would be distorted by projection. For planar surfaces, the distortion from projection can be approximated by a 3x3 projective transformation, known as a planar homography. A planar homography is a general class of transformations that includes rotations, translations, scaling, and arbitrary affine transformations as special cases. Planar homographies are invertible and are easily estimated from the observed position of any known 4 points on a surface. An inverse planar homography can be applied to an image before projection so that the image appears on a planar surface without deformation over a wide range of positions and orientations.

The environmental surfaces suitable for display use, and the required homographic transformation for each surface could, in principle, be determined from a detailed 3D model of an environment. However, the level of detail and precision necessary for such estimation would require a prohibitively expensive process of surveying and modeling the environment. Such 3D modeling can be avoided by directly measuring the required information using a camera.

Several camera-based methods of calibrating projector-display systems have been presented in [4, 11]. They are

based however on a time consuming 3D modeling of the screen, which is inappropriate for potentially fast moving displays. Moreover they use a calibrated camera to generate the model of the environment, which is an important constraint for the system infrastructure.

A homography-based automatic calibration method has been described in [6]. Though it uses a non-calibrated camera to observe the scene, it has a significant disadvantage of requiring that the camera is able to detect the screen boundaries. In this paper we present two display applications that include automatic image rectification. In the first application, an SVP and a non-calibrated remote camera enables dynamic image rectification on planar surfaces localized during an offline auto-calibration of the system (Section 3). In the second application, an SVP unit with its associated video camera enables dynamic image rectification on screens delineated in the projector field of view by their four boundaries (Section 4).

The possibility of rendering interfaces almost anywhere in an environment generates questions concerning how to control and interact with flexible displays. An interesting example of automatic control of a steerable display has been the subject of [3]. The display was automatically redirected to the most appropriate surface close to the user. In this paper we address the problem of enabling the user to control the projection direction. We present a tool that can be used as a means to transport a display from one surface to another (Section 4). We describe a robust tracking method that enables us to create a portable screen surface. Robust tracking of the screen boundary enables the planar homographic rectification to be re-computed dynamically as the display is moved.

## 2. Camera-Projector systems

Even for a rigidly mounted projector, a certain degree of steerability can be achieved in software, by restricting the size of the image to be projected to a subset of the available pixel array of the projector. For example, in [9] an Euclidean transformation was used to enable a projection of a circular interface to adapt its orientation with respect to the position of the users around a round table. The image may be electronically "steered" within the projector's pixel array, by modifying the transformation from image to projector.

Extending the display surface requires augmenting the range of angles to which the projector beam may be directed. A low cost technique for redirecting the projection beam with the use of a motorized mirror has been described in [4]. Their technique uses off-the-shelf components and has a capability to project very bright images within a cone restricted by the mirror-projector configuration. A prototype of a device having much less spatial constraints has




Figure 1: Classification of the calibration methods

been presented in [2]. In this device, a steerable pan-tilt platform was constructed to carry a small video-projector and a camera. With a trend towards increasingly small video projectors and cameras, this approach will become increasingly attractive.

In [11] projector-screen calibration methods are classified as passive vs. active, and online vs. off-line. We propose to add to this classification the number of significant dimensions of the acquired scene model (Figure 1). Off-line calibration methods have been used for static applications, where the projector-screen configuration was assumed to be constant in time and excluded eventual transitions between different configurations. Such an approach has been presented in [3, 4], where each projection surface was predefined and calibrated as a planar screen by interactively defining an appropriate homography. One might say that in this case the environment was modeled locally as a 2D screen surface. An automatic 2D model acquisition has been presented in [6], where the homography matrix was determined by an image processing method based on an uncalibrated camera. Though easy to apply, this approach relies on a strong assumption that the screen surface is planar, uniform in color and has distinguishable borders. Such constraints are often violated when projecting on unconstrained surfaces in real environments.

For most static or semi-static applications, where the screen is assumed to be planar, a 2D scene model of the environment is sufficient. Three dimensional environmental models can be used to enable rendering of images on screens with complex geometries. However, most such systems generate correct views when seen only from a single point in the scene. For example, in [5] imperceptible structured light is used to acquire the 3D geometry model of the screen in an augmented office application. Using some of the projectors energy to generate the structured light patterns visibly decreases the image contrast. An alternative method based on image processing techniques, in particular on image feature detection, has been shown in [11]. Both of these systems assume a calibrated projector-camera system.

### 3. The steerable projector

Our system has been designed to enable experiments in which a room is used as an interactive display surface. In particular, our system must permit experiments in which interactive display surfaces are automatically adapted to the relative positions and orientations of users. To meet this challenge, we have constructed a Steerable Video Projector (SVP) platform that provides a video-projector with two mechanical degrees of freedom: pan and tilt.

The SVP is composed of an XGA-resolution video-projector, a computer controlled pan-tilt platform and a CCD, PAL video camera. The mechanical configuration is shown in Figure 2. The SVP mechanical performance is




Figure 2: Scheme of SVP

presented in Table 1. In order to reuse the communication protocol and hardware interface from the pan-tilt cameras installed in our Augmented Meeting Environment laboratory we constructed our system by cannibalizing a widely used commercially available pan-tilt camera. We note that this environment also includes steerable cameras mounted in each of the four upper corners of the room, close to the ceiling. The serial computer interface for the pan-tilt head was connected to high performance stepper motors, enabling rotation of a load of up to approximately 2.5kg. The camera and lens assembly was detached from its pan-tilt base and rigidly fixed to the projector mounting. A mechanical assembly was designed so that the SVP could be mounted on the ceiling in an office-like augmented environment.

Table 1: Rotation platform mechanical performance

	Pan	Tilt
Rotation range	$\pm 177^\circ$	$+90^\circ$
Angular resolution	$0.11^\circ$	$0.18^\circ$
Angular velocity	$146 \frac{deg}{s}$	$80 \frac{deg}{s}$
Response time	$\sim 2ms$	$\sim 3ms$

### 3.1. Rectifying projected images

Our approach to image geometry correction is based on the observation that in most office or room-like environments the majority of surfaces suitable for use as projection screens are planar. We present a two step method for oblique distortion correction in images projected on any part of arbitrarily oriented planar surfaces. The method requires the use of a motorized video-projector (with known optical parameters), and an arbitrary placed non-calibrated video camera. The camera must see the entire projected images.

First the camera acquires necessary data to verify the shape, size and orientation of the screen surface (Section 3.2). Second for each projector-screen configuration the system calculates the necessary pre-warp that has to be applied to the projected image in order to compensate the oblique distortions (Section 3.3). This method can be extended by a third step correcting distortions that are due to a non-planar shape of the screen. It could also be based on algorithms presented in [1] or [11].

### 3.2. Finding the screen

The first phase of our method is an automatic off-line calibration based on the observation that a perspective projection transforms a line into another line. The projection of an image onto a planar surface, as well the observation of planar surface with a camera, are described by a simplified (or degenerate) form of perspective transformations.

Since a composition of two perspective transformations conserves straight lines, this feature can be exploited to provide a simple method to determine surfaces that are suitably close to planar to be used as a display surface for a rectified image. A grid pattern is projected by the video projector and captured by any available video cameras. Using algorithms described below in Section 4 we find the shape of the pattern in the camera image. Any discontinuity of the projected lines indicates a corner-like distortion of the screen. A curved appearance of the lines in the pattern indicates that the surface is not planar.

The system records projector orientation angles corresponding to each detected surface. As a result, a two dimensional model of the environment is created in which each surface has a representation as a region.

Conservation of lines is also used to establish the configurations in which the optical axis of the video-projector is parallel to the normal of the screen surface. For this, two patterns of two dots are projected and moved across the surface within the sight of a camera. The dots lie on one of two lines (a vertical and horizontal with respect to the image borders) crossed in the optical center of the image. Each dot is color coded in order to identify it in each camera image.

A vertical aligned dot pattern is projected when estab-

lishing the angle between the screen normal and the optical axis of the projector in the vertical plane. While rotating the projector around its vertical axis, the points are detected in the camera image. Each of the points follows a trajectory which is a function of the projector-screen-camera configuration. A linear regression algorithm is applied for each sequence of colored points seen by the camera and a mean point-to-line distance is calculated. The minimum distance corresponds to a trajectory aligned with the intersection of the screen with a perpendicular plane passing through the focal point of the projector. One of the projector's orientation angles can be measured with respect to the screen by detecting this minimum and the corresponding point coordinates in the projected image. In a similar manner, the second orientation angle is found using a color-pattern of horizontally aligned dots.

The ability of the automatic environment modeling technique to determine the orientation of planar surfaces is limited to certain projector-screen configurations. In particular, with the presented technique the system is unable to find the orientation of small-area surfaces. Thus the system is unable to correct image distortions using the method described below (Section 3.3). Therefore in Section 4 we present an alternative projector-screen calibration technique that enables us to project rectified images also on small-size screens.

### 3.3. Projector-screen distortions model

The goal of the distortion correction algorithm is to determine a transformation  $M$  from source image  $X_{si}$  to a projected image  $X_{pi}$ .

$$X_{pi} = MX_{si}$$

This transformation should be such that after projection, the projected image appears as if the source image had been projected to a perpendicular, planar screen, up to a scale change  $S$ .

$$SP_0X_{si} = PX_{pi}$$

We define the origin of the projected image frame  $R_{pi}$  and the source image frame  $R_{si}$  as the intersection of the projector optical axis with the image. The origin of a world reference frame  $R_W$  is in the focal point of the video-projector optics. The  $x_w$  axis is aligned with the vertical rotation axis of the projector. The  $z_w$  axis points to one of the planar surfaces in the environment and is parallel to the normal of the surface (Figure 3). The two angles  $\psi$  and  $\theta$  defining the position of the SVP equal zero when the projector optical axis is aligned with the  $z_w$  axis of the world frame. Therefore the projection  $P$  of a point in the projected image frame  $R_{pi}$  can be decomposed as:

$$X_S = PX_{pi} = P_pCX_{pi}$$




Figure 3: Projection model

Where the matrix  $C$  corresponds to a transform from the projected image frame  $R_{pi}$  to the world frame  $R_W$ . It has the following form:

$$\begin{bmatrix} x_W \\ y_W \\ z_W \\ w \end{bmatrix} = \begin{bmatrix} R(\psi, \theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{pi} \\ y_{pi} \\ 0 \\ 1 \end{bmatrix}$$

Where  $f$  is the focal length of the projector optics and  $R$  is the projector rotation matrix. The projector projection matrix  $P_p$  is composed of the distance  $d$  of the screen to the focal center of the projector:

$$P_p = \begin{bmatrix} d & 0 & 0 & 0 \\ 0 & d & 0 & 0 \\ 0 & 0 & d & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

The unknown distance  $d$  can be arbitrary chosen, since this results only in a change of scale. One can say that transformations  $C$  and  $P_p$  project points from the projected image frame onto a virtual screen at a distance  $d$  from the projector focal center. The boundaries of the projection are determined such that the corners of the projected image coincide with the corners of the virtual screen.

Within these boundaries the system can display a rectangle (or any other figure)  $X_{Sr}$  that users would like to see on the real screen. The rectangle size can be adapted so that apparent size is independent of the orientation or position of the screen surface by calculating the transformation from the four screen corners.

In order to determine the new coordinates of a rectified image  $X_{Sr}$  one can translate the image  $X_{S0}$  that would result from a perpendicular projection  $P_0$  by a translation  $T$ .

$$\begin{bmatrix} x_{Sr} \\ y_{Sr} \\ z_{Sr} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & d \tan \psi \\ 0 & 1 & 0 & d \tan \theta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{S0} \\ y_{S0} \\ z_{S0} \\ 1 \end{bmatrix}$$

Which can also be expressed as:

$$X_{Sr} = TP_0X_{si}$$

Finally the desired figure  $X_{Sr}$  can be back-projected to the  $R_{pi}$  (projected image frame).

$$X_{pi} = T_B P_B R^T X_{Sr}$$

Where  $P_0 = P_p C$  for  $\psi = 0$  and  $\theta = 0$ . The matrix  $R^T$  is the transpose in homogeneous coordinates of the rotation matrix  $R$ . Matrixes  $P_B$  and  $T_B$  assure the back-projection from the scene to  $X_{pi}$ .

$$P_B = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & f & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad T_B = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & -f \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Since we know the parameters of the projector and the orientation angles, it is sufficient to apply the mapping  $M$  to the desired image shape  $X_{si}$ .

$$X_{pi} = T_B P_B R^T T P_0 X_{si} = M X_{si}$$

The resulting image  $X_{pi}$  should not exceed the projector’s pixel array, otherwise we have to apply an appropriate scaling.

Thanks to this approach we are able to dynamically rectify the image while sweeping it over a planar surface without having to recalibrate the projector-screen system for each position of the SVP.

## 4. Application: a novel user-interface

The association of a steerable video projector with a camera enables to convert an ordinary room into an “interactive environment”. If the correction matrix for a new display surface is based on previously created environment model (Section 3.2) or a precalibration method, as proposed in [4], then interaction is limited to pre-defined static surfaces. Dynamic recalibration allows us to extend interaction to mobile, hand-held surfaces.

In this section we illustrate the use of the steerable projector by describing a new form of interactive device using only a small hand-held screen. We call this a “Portable Display Screen” or PDS. The steerable video-projector with its associated camera will be used to transfer an interactive display from a large screen onto the hand-held screen, automatically correcting for 3D translations and rotations as the screen is moved. The SVP camera can be used to detect when the user places his fingers or hand over a part of the screen in order to select or manipulate projected information. Additionally the PDS will be used as a means of transferring the projected information between screen surfaces in the environment.

A typical application scenario for the PDS would be as follows: at the end of a presentation where slides have been displayed on a wide screen, a user may wish to move the display from a main wide screen somewhere else. Introducing the hand-held screen over the projected image indicates that user wishes to move an interactive display onto the hand-held screen. Our steerable projector-camera pair allow to maintain the image on the small hand-held screen. As the PDS is removed from the previous location of the wide screen, its corners are tracked to allow the projector to be steered in order to maintain the image on the PDS. Tracking the screen corners enables also the projected image to be rectified. The information can be transferred from the PDS on another display surface by placing the PDS over it. By doing it the user can define a new display surface in the environment. By knowing the position of the four corners of the PDS, the system immediately knows the mapping between the projected image and the newly chosen planar surface.

### 4.1. The PDS

The PDS can be made from any flat screen that is small enough to be held in a hand. The screen should have a uniform, light-colored surface delineated by a clearly marked border. An example is shown in Figure 4. It is a piece of white cardboard with black boarders.



Figure 4: A hand-held screen used to construct a portable display surface

### 4.2. Initialisation of the user-interface application

In order to initialize our PDS demonstration, we added a small red rectangle to the upper right region of the projected image (Figure 5). Passing the hand-held screen over this red rectangle initializes the PDS.

### 4.3. Tracking the PDS

Both tracking and image rectification are based on the position of the four corners of the hand held screen. However,



Figure 5: Detection of a rectangle in the upper right part of the board

the corners, themselves may be occluded. Rather than try to directly detect the corners, we use a Hough transform to estimate and track the four screen boundaries. The corners are then determined as the intersection of the four dominant edge-lines subject to constraints on their difference in position and orientation.

The PDS is assumed to be held in a hand. Thus some part of the display is occluded by a skin colored region. Prior to boundary detection, a histogram based skin color detector is applied to remove skin colored regions. Edge points are then detected using the gradient magnitude computed by convolution with a gradient of a Gaussian at an appropriate scale,  $\sigma$ .

$$\text{Gradient}(\sigma, x, y) = \sqrt{L_x(\sigma, x, y)^2 + L_y(\sigma, x, y)^2}$$

where  $L_x$  and  $L_y$  the first Gaussian derivatives. Points of high contrast are used by the Hough transform to estimate the screen boundaries, as shown in (Figure 6). We are currently experimenting with the use of a Laplacian of Gaussian operator that has been tuned to the width of the border. A ridge detector is applied to the Laplacian of Gaussian image, and only points along a sufficiently long, straight ridge are used to compute the Hough transform.

The event which represents the detection of a clear rectangle in the red region of detection should be determined with a very high accuracy and detection rate. That is why a more complex module which additionally computes robust ridges with maxima of Laplacian [7] can be used as confirmation: Figure 7.

Screen boundaries are selected as maxima in the Hough space and represented in Figure 8 with red crosses. The line position is defined by the following equation:  $\rho = x \cos \theta + y \sin \theta$  with  $\theta \in [0; \pi]$ . Hough transforms have relatively large cost in memory and computational time. Nevertheless, lines are only detected in regions of interest and not in the whole image, thus speeding up the computation




Figure 6: Enhanced image of the gradient (threshold and remove of skin color pixels)




Figure 7: Robust ridges directly computed on the source image (without enhancement). Red lines are positive Laplacian values and blue are negative ones.

and reducing the memory requirements. Moreover,  $\theta$  can be incremented in intervals of 2 or more.

#### 4.4. Robust tracking of 4 segments

Once a rectangle has been clearly detected in the region of interest, a robust tracking algorithm is initialized. The tracker follows targets as points in the Hough space with adaptive regions of interest for each set of four local maxima in that space. The mobile camera and video-projector are steered together in pan and tilt so as to maintain the center of gravity of the four points in the center of the projected image.

The aim is to continuously track 4 local maxima in the Hough space, predict their new locations, re-compute the Hough space for the new image and test among the new local maxima the best four which define the new rectangle location. Thus, segments are first independently tracked giving for each segment a list of possible new locations (Figure 10). Note that in order to speed up the algorithm the Hough space is only computed for angles  $\theta$  that belong to the local maxima regions of interest.




Figure 8: Hough space with the fifteen highest local maxima represented by red cross




Figure 9: List of lines detected using the local maxima detected in the Hough space (Figure 8)

The list of new lines is sorted by their confidence computed from a-priori constraints such as the position and slope of the previous segment (i.e  $\rho$  and  $\theta$ ). These confidence factors are computed using a Mahalanobis distance:

$$\text{ArgMin}_{Y_k} (Y_k - X^*)^T (C + C^*)^{-1} (Y_k - X^*) \quad (1)$$

with:

- $k$  the index of the lines found.
- $Y_k = (\rho, \theta)$  the characteristics of the segment.
- $C = (\sigma_\rho, \sigma_\theta)$  the assumed error position.
- $X^*$  the prediction of the new position.
- $C^*$  the prediction of the assumed error position.

An adaptation of the RANSAC algorithm is used to apply strong constraints to all possible lines (Figure 9). The result selects the four most likely lines defining the new segments position. These strong constraints ensure that segments are roughly pairwise parallel and that the distance between 2 parallel segments are similar to distance observed in the previous image.




Figure 10: For one of the four maxima tracked: list of hypothetical new local maxima compared to its assumed position given by the prediction

The robust tracking algorithm runs as follows:

1. Compute the enhanced Gradient image for the predicted region of interest.
2. Threshold the resulting image and remove pixels that are too far from previous lines positions.
3. Find maxima in selected regions of the Hough space and for each segment of the previous rectangle location: list and sort the matching lines.
4. Apply geometric constraints to the possible lines adjustments and select the 4 new maxima.
5. Update new characteristics and predict the new positions of each 4 local maxima in the Hough space defining four new regions of interest in the Hough space.

A Kalman filter is applied to the region of interest where the PDS was found. If no rectangle in this region is found, the size of the region of interest is increased by a factor of 1.25. A confidence estimate is maintained, and if confidence falls below a threshold, the tracker is halted and the initialization procedure is re-started.

#### 4.5. Projection of the image onto the board

The relation between a pixel of the projector image with homogeneous coordinates  $[x_1 w, y_1 w, w]^T$  and its corresponding point into the PDS  $[u_1, v_1, 1]^T$  is a projective transform with eight degrees of freedom. This transformation may be expressed in homogeneous coordinates as a 3 by 3 projective transformation,  $H$ . The nine coefficients of  $H$  can be exactly determined by the observed position of 4 points combined with the constraint that  $|\vec{H}| = 1$ . These four points must be non-collinear. That is, no subset of three

points may lie on the same line in the observed image. In this case,  $H$  is a homographic projection that preserves the mapping from one image to another, and is thus invertible. The four corners of the display screen, tracked by the Hough transform, provide exactly such a set of four points.

Observing the four corner points of the display screen provides the homographic transformation,  $H_s^c$  from the screen to the camera. As shown in [6] image rectification requires the homography from the screen to the projector  $H_s^p$ .

$$H_s^p = H_c^p H_s^c$$

Since the camera is rigidly mounted on the projector, the transformation  $H_c^p$  is roughly constant. It can be calibrated during system setup by interactively aligning the boundaries of a projected rectangle with the boundaries of the rectangular screen. Thus the observed homography to the camera can be converted to the required homography from the screen to the projector.



Figure 11: projection of the image into the PDS

Figure 11 shows the result of a projection by the video-projector of a image onto a hand held screen. This projection should leave the PDS borders clear in order to continue to be easily tracked.

## 5. Conclusion

This paper describes a system in which a camera is rigidly mounted to a video projector, and the camera-projector pair are steered together. The camera may be used to track a mobile screen both for steering the camera-projector assembly, and to estimate a homographic transformation so that the projected image arrives at the screen in rectified form. A method for automatic environment model capture is introduced. The method takes advantage of the projectors ability to control the position of the projected images in the environment. It enables the system to keep projecting on a planar screen an image with a constant aspect ratio even while rotating the projector. This paper has also presented a new

interaction device in a form of a portable screen surface. A technique is described to redirect the video-projector to a new display location by transporting the screen on the portable display surface.

## Acknowledgments

This work has been partially funded by European IST projects FAME and DETECT.

## References

- [1] R. Raskar et al. iLamps: Geometrically aware and self-configuring projectors. In *Appears ACM SIGGRAPH 2003 Conference Proceedings*.
- [2] N. Nakamura and R. Hiraïke. Active projector: Image correction for moving image over uneven screens. In *Companion of the 15th Annual ACM Symposium on User Interface Software and Technology*, pages 1–2, October 2002.
- [3] G. Pingali, C. Pinhanez, T. Levas, R. Kjeldsen, and M. Podlaseck. User-following displays. In *Proceedings of the IEEE International Conference on Multimedia and Expo 2002 (ICME'02)*, 2002.
- [4] C. Pinhanez. The everywhere displays projector: A device to create ubiquitous graphical interfaces. In *Proceedings of Ubiquitous Computing 2001 Conference*, September 2001.
- [5] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays. In *Proceedings of the ACM SIGGRAPH'98 Conference*.
- [6] R. Sukthankar, R. Stockton, and M. Mullin. Smarter presentations: Exploiting homography in camera-projector systems. In *Proceedings of International Conference on Computer Vision*, 2001.
- [7] H. Tran and A. Lux. Extraction de caractéristiques locales: Blobs et cretes. RIVF Hanoi, Viet Nam, Feb 2003.
- [8] J. Underkoffler and B. Ullmer and H. Ishii. Emancipated pixels: Real-world graphics in the luminous room. In *Proceedings of ACM SIGGRAPH*, pages 385–392, 1999.
- [9] F. Vernier, N. Lesh, and C. Shen. Visualization techniques for circular tabletop interfaces. In *Advanced Visual Interfaces*, 2002.
- [10] P. Wellner, W. Mackay, and R. Gold. Computer-augmented environments: back to the real world. *Communications of the ACM*, 1993.
- [11] R. Yang and G. Welch. Automatic and continuous projector display surface calibration using every-day imagery. In *CECG'01*.