

Automatic Keystone Correction for Camera-assisted Presentation Interfaces

Rahul Sukthankar^{1,2}, Robert G. Stockton¹, and Matthew D. Mullin^{1,*}

¹ Just Research, 4616 Henry Street, Pittsburgh, PA 15213, U.S.A.
{rahuls,rgs,mdm}@justresearch.com

² The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.
rahuls@cs.cmu.edu
<http://www.cs.cmu.edu/~rahuls>

Abstract. Projection systems have become the ubiquitous infrastructure for presentation technology. However, unless the projector is precisely aligned to the presentation screen, the resulting image suffers from perspective (keystone) distortions requiring manual optical or digital correction. This tedious process must be repeated whenever the projector or screen is moved and is increasingly relevant given the emerging trend towards highly-portable LCD projection systems. This paper presents a presentation interface that pre-warps the image to be projected in such a way that the distortions induced by the projector-screen geometry precisely negate the warping. An uncalibrated, low-resolution digital camera is used to infer the projector-screen geometry and to automatically determine the pre-warping parameters. This vision-based system is augmented with a natural interface that enables the user to interactively refine the suggested rectification. Arbitrary distortions due to projector placement are negated, allowing the projector (and camera) to be placed *anywhere* in the presentation room — for instance, at the side rather than the center of the room. Our solution works with existing projector hardware, and could easily be incorporated into the next generation of LCD projector systems.

Introduction

Projection systems are the ubiquitous infrastructure for presentation technology, and portable LCD projectors have become increasingly popular for multimedia presentations. However, unless the projector is carefully aligned to the projection surface (screen), the resulting image on the screen appears distorted, or *keystoned*¹. Keystoning is undesirable, not only because viewers find this warping to be distracting, but also because the distortion is detrimental to the interpretation of visual information such as graphs, bar-charts and technical drawings.

* Rahul Sukthankar is now with Carnegie Mellon University and Compaq Cambridge Research Lab. Robert Stockton and Matthew Mullin are now with WhizBang! Labs.

¹ In related work, “keystoning” refers specifically to a symmetric, trapezoidal distortion caused by projector pitch misalignment. Here, the term refers to the broader class of distortions caused by *any* misalignments in projector position or orientation.

Keystoning can be prevented by aligning the projection system’s optical axis so that it is perpendicular to the screen, and ensuring that the image is not rotated with respect to the screen. For fixed projectors that can be mounted from the ceiling and carefully aligned once, these constraints are surmountable; however, portable projectors require alignment at the start of each presentation session. This manual process tedious; it can also be impractical to align a portable projector in a manner that eliminates all keystoning effects since optimal alignment may place the projector in an awkward position (such as in the middle of the audience). This motivates the need for a better presentation interface: one that allows arbitrary placement of the projector.

Sophisticated projectors now offer *manual* keystone correction that allow the user to counter the limited class of distortions caused by *vertical* misalignment. Unfortunately, this form of keystone correction only addresses projector pitch² and cannot compensate for distortions due to projector roll or yaw. More importantly, requiring the user to directly adjust the projector optics is undesirable from a user-interface standpoint: precise correction requires tweaking several coupled parameters, whose effect on the projected image is non-intuitive.

This paper presents a fully-automatic method for keystone correction. The two key concepts are: (1) a digital camera is used to observe the projected image; (2) the image to be displayed is pre-warped so that the distortions induced by the projection system will exactly undo the keystone distortion. The result is that an arbitrarily mounted projector (in an unknown orientation) displays a perfectly aligned and rectilinear image (see Figure 1).

Vision-based Keystone Correction

The hardware requirements are modest: a standard computer, a low-resolution digital camera (e.g., Logitech QuickCam), and a standard LCD projector. The projector can be mounted anywhere, as long as the image falls entirely within the projection screen area (e.g., the projector could be placed near the edge of the room, even at an angle). The camera must be mounted such that the projection screen is within its field of view (but the camera can be placed anywhere in the room, and need not be level).

Our method for keystone correction is summarized as follows. (1) Determine the mapping between points in the computer display and the corresponding points in the camera image. (2) Identify the quadrilateral corresponding to the boundaries of the projection screen in the camera image. From this, compute a possible mapping between the projection screen and the camera image frame. (3) Infer a possible mapping from the computer display (source image frame) to the projection screen based upon the mappings computed in the previous two steps. (4) Determine an optimal placement for the corrected image on the projection screen. This is the largest rectangle that is completely contained within the projection of the computer display (i.e., the keystone quadrilateral in Figure 1).

² Furthermore, current commercial keystone-correction systems are restricted to rectifying small misalignments in pitch, typically no more than $\pm 12^\circ$.




Fig. 1. An illustration of the apparatus. The computer is connected to a projector, and the projected image is observed by the camera. The positions, orientations and optical parameters of camera and projector are unknown. Due to projector misalignment, the rectangular screen appears as a distorted quadrilateral (shown shaded). However, by properly pre-warping the source image, the projected image appears rectilinear (shown by the white rectangle enclosed by the keystone quadrilateral). The pre-warping parameters are automatically determined by the projector-camera system calibration.

(5) Pre-warp each application image to correct for keystone. These steps are described below in greater detail, and results are presented in Figure 2.

Projector-Camera System Calibration

The goal of this step is to determine a mapping between points in the source image and the corresponding points in the camera image (see Figure 3). Surprisingly, although the positions, orientations and optical parameters of the camera and projector are unknown, this mapping (referred to as T in the remainder of this paper) can still be inferred.

First, we note that the mappings from the source image frame to the projected image frame, and from the projected image frame to the camera image frame are each perspective transforms. When these two transforms are composed (i.e., the projection of the source image is viewed through the camera), the resulting mapping, while not necessarily a perspective transform, can be expressed as a projective transform:

$$(x, y) = \left(\frac{p_1 X + p_2 Y + p_3}{p_7 X + p_8 Y + p_9}, \frac{p_4 X + p_5 Y + p_6}{p_7 X + p_8 Y + p_9} \right),$$

where (x, y) is a point in the source image frame, (X, Y) is the corresponding point in the camera image frame and the parameters $p_1 \dots p_9$ are the unknowns




Fig. 2. A rectangular source image (top left) appears distorted when projected from an off-center projector (middle & bottom left). Using our method, the source image is pre-warped (top right). The resulting projected image is rectilinear and perfectly aligned to the projection screen (middle & bottom right). The middle row of images was captured with the low-quality camera used in our apparatus, and the bottom row was captured with a high-quality camera to illustrate how the presentation would appear to an audience member in the room.

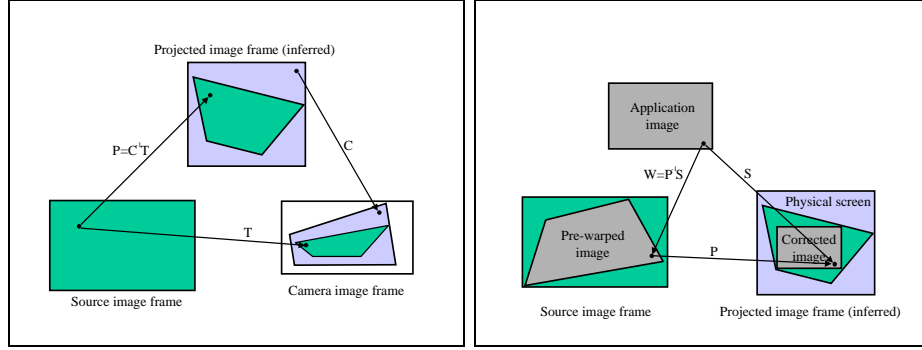


Fig. 3. Left: the relationships between the three frames of reference corresponding to the computer display (source image frame), camera (camera image frame) and physical screen (projected image frame). Note that the last can only be indirectly observed by our system, through the camera. T is obtained using the calibration method; C is obtained by locating the screen within the camera image. Finally, P , the mapping responsible for the keystone distortion, is derived mathematically as T composed with the inverse of C , or $P = C^{-1}T$. **Right:** the application image can be appropriately distorted (pre-warped), using the mapping W so that it appears rectilinear after projection through a misaligned projector (modeled by the mapping P).

to be determined. Although there are 9 unknowns in this equation, there are only 8 degrees of freedom ($\sum_i p_i = 1$). Four point correspondences (where each point provides two constraints) are therefore necessary. Fortunately, in our system, these point correspondences can be automatically obtained by projecting a known rectangle into the environment, and observing the locations of its corners through the camera. Given four points, a unique solution for the parameters is obtained using standard linear algebra techniques.³

Identification of Projector Screen Boundaries

The general problem of keystone correction requires pre-warping a source image such that, when projected, its edges will be aligned parallel to the boundaries of the projection screen. In the typical case, where the projection screen is rectangular, this implies that the corrected image will also be rectangular. More precisely, the edges of the corrected image should converge towards the same vanishing points as the corresponding edges of the physical screen.

To do this, we must accurately locate the boundaries of the physical projection screen, which is a challenging computer-vision task. Since presentation rooms are typically darkened, the system first projects a bright white image onto the screen. By assuming that the projection surface is uniformly light in color, it is possible to extract the boundaries of the projection screen, even outside the boundaries of the projected image. This is done by a region-growing

³ When greater than four points are available, a least-squares solution is used.

procedure where adjacent pixels with similar intensities are grouped into connected components, and all such components near the bright projected image are merged. This can be seen in Figure 3 (left): the outer quadrilateral in the camera image frame corresponds to the projection screen while the inner quadrilateral corresponds to the bright projected image. Although the camera images are low-resolution, high calibration accuracies are obtained by fitting lines to the edges of the quadrilateral, and computing intersections (corresponding to the coordinates of the corners of the screen) to sub-pixel accuracy. From these corners, and an assumption that the physical screen is rectangular, we derive a mapping C that transforms points from the projected image frame to the camera image frame (within a scale factor).

Inferring mapping P

As shown in Figure 3 (left), the mapping (P) from the source image frame to the projected image frame may be computed from the two known mappings T and C derived above, by composing T with the inverse of C : $P = C^{-1}T$. In our system, P models the configuration of the projector setup that is directly responsible for the keystone distortion. Our task is therefore to pre-warp the application image, so that it is rectilinear (distortion-free) after being mapped through P .

Determining optimal placement for corrected image

The computer display (source image frame) projects to an arbitrary quadrilateral (projected computer display) in the projected image frame as shown in Figure 3 (right). Since the pre-warped image can only be displayed within the bounds of the computer display, the corrected image must lie within the bounds of this quadrilateral, and (for best viewing) should be as large as possible. This is equivalent to finding the largest rectangle with appropriate aspect ratio within the projected computer display; a heuristic optimization is used to compute the dimensions of this rectangle.

Now, given the desired size and location of the corrected image, we compute a mapping S that scales and shifts the application image to lie within this area. We can compute the necessary pre-warping transformation W from the two known transformations S and P by the formula $W = P^{-1}S$, as shown in Figure 3 (right).

Interactive refinement

The transformation, W , computed above is a candidate model for the correction required to undo projector misalignment. However, since W may be inaccurate due to errors in the calibration process, the user is given the opportunity to interactively refine the warping parameters. The rectangle bounding the source image is warped using the candidate W and projected onto the screen. If the

calibration process were successful, the projected image would appear to be a rectangle aligned with the screen. The user, through a GUI, can directly manipulate the corners of the rectangle in order to fine-tune the alignment; this feedback is used by the system to refine W . Users of the presentation system have found this interface to be far superior to directly adjusting the eight degrees of freedom in W (e.g., using slider bars) since it allows the same level of control (the four corners of the rectangle provide eight parameters, uniquely constraining W) while being very intuitive. In fact, several users report that this interface alone (without the camera) enables them to perform excellent manual keystone-correction, indicating that such an interface would significantly improve current (manual) keystone-correction projectors.

Image warping for keystone correction

Once a correct W has been determined, the pre-warped source image is created as follows. For each pixel in the pre-warped image, we find the corresponding point in the application image by applying the inverse mapping, W^{-1} . Since the computed point is real-valued, it will not typically correspond to a single (integer) pixel in the application image. Therefore, the four pixels in the application image that are closest to this point are blended using bilinear interpolation [1]. Note that the image warping computations are well-suited to exploit the accelerated 3-D graphics hardware available in typical computers. The application image is embedded in a virtual region of black pixels to ensure that only pixels within the corrected image are illuminated. The resulting image, when projected through the misaligned projector, appears undistorted.

Conclusion

Our system provides a better interface for presentation systems by allowing users to place portable projectors anywhere in the room, and automatically rectifying the resulting distortions. In particular, projectors may be placed at the side of the room, where they will not interfere with the audience. Recent related work in the area of teleconferencing [6] also demonstrates the benefits of distortion-correcting projector-based interfaces.

The availability of low-cost cameras and advances in processing power are making perceptual interfaces accessible to a general audience. Our automatic keystone correction system is one component in a larger *camera-assisted presentation interface*, prototyped at Just Research. Unlike typical perceptual interfaces, the camera-assisted presentation interface observes the presentation rather than the user, becoming a member of the audience. The user's interactions with the audience (e.g., laser pointer gestures) are observed and interpreted by the camera-assisted presentation system, allowing the user to control the presentation in a natural manner.

Although the system as described in this paper works with existing projector hardware, it could easily take advantage of newer projector technology. For

instance, if the projector's position, location or optical characteristics could be computer-controlled, then the mapping W could be translated into commands to adjust these directly. Similarly, by integrating the camera hardware into the projector (e.g., the Proxima Cyclops [5]), the calibration problem is significantly simplified (T is the identity). We anticipate that newer generations of projectors will incorporate the automatic vision-based keystone-correction techniques described here.

If sufficient computational resources are available on the user's computer *during* the course of a presentation, then the keystone correction system could perform automatic online recalibration. This would have important benefits. For instance, if the camera or projector were accidentally moved during the course of the presentation, the system would automatically detect the disturbance by observing the discrepancy between expected and observed camera images. The pre-warp transformation would then be automatically adjusted so that the projected image could remain steady in spite of the motion. In an extreme scenario, one could even envision mounting the camera and/or projector on a moving platform (e.g., such as the user's body) with no adverse effect in the projected image.

The automatic keystone correction system described in this paper has been combined with other vision-based presentation interfaces (such as laser pointer presentation control) to create a *camera-assisted presentation interface*. The complete system has been deployed since December 1999, receiving very favorable response and attracting considerable commercial interest.

Acknowledgments

Thanks to Rich Caruana and Terence Sim, with whom the initial ideas for vision-based keystone correction were discussed, and to Gita Sukthankar for valuable feedback on the paper. Provisional patent applications for the inventions stemming from this work have been filed by Just Research [2, 3, 4].

References

- [1] J. Foley, A. van Dam, S. Feiner, and J. Hughes. *Computer Graphics: Principles and Practice*. Addison Wesley, 1993.
- [2] M. Mullin, R. Sukthankar, and R. Stockton. Calibration method for projector-camera system. U.S. Patent Filing, 2000. Provisional patent filed in 1999.
- [3] R. Sukthankar, R. Stockton, and M. Mullin. Automatic keystone correction. U.S. Patent Filing, 2000. Provisional patent filed in 1999.
- [4] R. Sukthankar, R. Stockton, M. Mullin, and M. Kantrowitz. Vision-based coupling between pointer actions and projected images. U.S. Patent Filing, 2000. Provisional patent filed in 1999.
- [5] G. Van Horn. Proxima's new Ovation+ projection panels do up multimedia. *Byte*, January 1995. <<http://www.byte.com/art/9501/sec12/art9.htm>>.
- [6] R. Yang, M. Brown, B. Seales, and H. Fuchs. Geometrically correct imagery for teleconferencing. In *Proceedings of ACM Multimedia*, 1999.