Detecting Photographic Composites of People

Micah K. Johnson and Hany Farid

Dartmouth College
Hanover, NH 03755, USA
{kimo,farid}@cs.dartmouth.edu
http://www.cs.dartmouth.edu/~{kimo,farid}

Abstract. The compositing of two or more people into a single image is a common form of manipulation. We describe how such composites can be detected by estimating a camera's intrinsic parameters from the image of a person's eyes. Differences in these parameters across the image are used as evidence of tampering.

Key words: Digital Tampering, Digital Forensics

1 Introduction

From the creation of tabloid covers to political advertisements to family portraits, the compositing of two or more people into a single image is a common form of manipulation. Shown on the right, for example, is a composite of actress Marilyn Monroe (1926-1962) and President Abraham Lincoln (1809-1865).

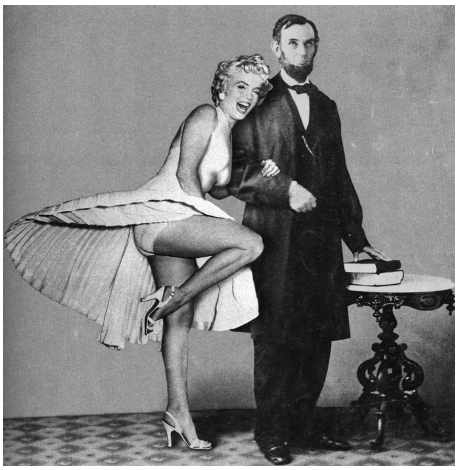

Composite of Marilyn Monroe and Abraham Lincoln (*by Jack Harris*).

Over the past few years the field of digital forensics has emerged to detect various forms of tampering. Forensic techniques have been developed for detecting cloning [2, 13]; splicing [11]; re-sampling artifacts [1, 14]; color filter array aberrations [15]; disturbances of a camera's sensor noise pattern [10]; and lighting inconsistencies [4, 6, 7]. Here we describe a new forensic technique specifically designed to detect composites of people. This approach estimates a camera's principal point from the image of a person's eyes. Inconsistencies in the principal point are then used as evidence of tampering.

In authentic images, the principal point is near the center of the image. When a person is translated in the image as part of creating a composite, the principal point is moved proportionally. Differences in the estimated principal point across the image can therefore be used as evidence of tampering. Shown in

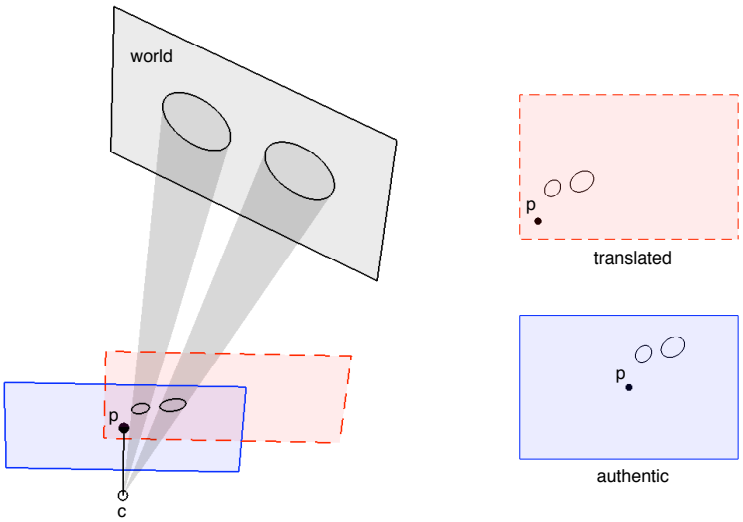

Fig. 1. Shown on the left is the projection of two circles (eyes) from a world plane onto an image plane (blue, solid line), with camera center c and principal point p. Also shown is a translated image plane (red, dashed line) which is equivalent to translating the image of the eyes in the original image plane (as shown on the right). Note that the translation results in the movement of the principal point.

Fig. 1, for example, is the projection of two circles (eyes) from a world plane onto an image plane. The point c denotes the camera center and the point p denotes the principal point (the projection of c onto the image plane). Also shown in this figure is a translated image plane which is equivalent to translating the image of the eyes in the original image plane. Note that the translation results in the movement of the principal point away from the image center.

We describe how to estimate a camera's principal point from the image of a pair of eyes. We then show how translation in the image plane is equivalent to a shift of the principal point. Inconsistencies in the principal point across an image are used as evidence of tampering. We show the efficacy of this approach on synthetic and real images and visually plausible forgeries.

2 Methods

In general, the mapping between points in 3-D world coordinates to 2-D image coordinates is described by the projective imaging equation:

$$\boldsymbol{x} = P\boldsymbol{X}, \tag{1}$$

where the matrix P is a 3×4 projective transform, the vector $\boldsymbol{X}$ represents a world point in homogeneous coordinates, and the vector $\boldsymbol{x}$ represents an image

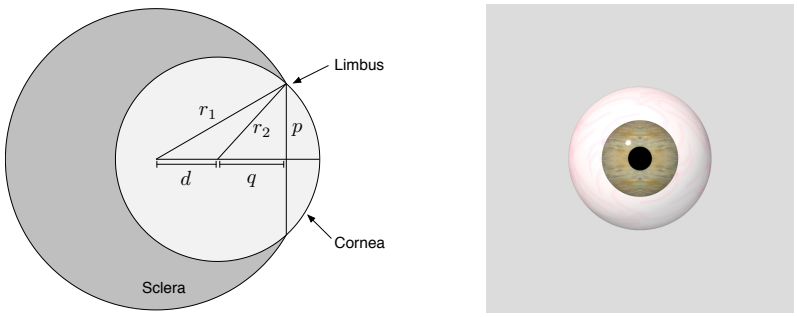

Fig. 2. A 3-D model of a human eye consisting of two spheres (left) and a synthetic eye rendered according to the model (right).

point in homogeneous coordinates. If all the world points $\boldsymbol{X}$ are coplanar, then the world coordinate system can be defined such that the points lie on the $Z = 0$ plane. In this case, the projective transformation P reduces to a 3×3 planar projective transform H, also known as a homography:

$$\boldsymbol{x} = H\boldsymbol{X}, \tag{2}$$

where the world points $\boldsymbol{X}$ and image points $\boldsymbol{x}$ are represented by 2-D homogeneous vectors.

We first describe how the homography H can be estimated from an image of a person's eyes and show how this transform can be factored into a product of matrices that embody the camera's intrinsic and extrinsic parameters. We then show how translation in the image plane can be detected from inconsistencies in the estimated camera's intrinsic parameters.

2.1 Homography Estimation

The homography H between points on a world plane and its projection on the image plane can be estimated if there is known geometry in the world: parallel lines, orthogonal lines, regular polygons, or circles [3, 9]. We will focus primarily on the known geometry of a pair of eyes (circles) to estimate the homography.

A simple 3-D model for an eye consists of two spheres [8]. The larger sphere, with radius $r_1 = 11.5$ mm, represents the sclera and the smaller sphere, with radius $r_2 = 7.8$ mm, represents the cornea, Fig. 2. The centers of the spheres are displaced by a distance $d = 4.7$ mm [8]. The limbus, the boundary between the iris and the sclera, is defined by the intersection of two spheres – a circle with radius $p = 5.8$ mm.

With the assumption that the two circular limbi are planar, the homography H, Equation (2), can be estimated from a single image of a pair of eyes. Intuitively, the limbi will be imaged as ellipses (except when the eyes are directly facing the camera) and the distortion of the ellipses away from circles will be

related to the pose and position of the eyes relative to the camera. We therefore seek the transform that aligns the image of the limbi to circles.

Points on the limbus in world coordinates satisfy the following implicit equation of a circle:

$$f(\boldsymbol{X};\boldsymbol{\alpha}) = (X_1 - C_1)^2 + (X_2 - C_2)^2 - r^2 = 0, \tag{3}$$

where $\boldsymbol{\alpha} = (\ C_1\ \ C_2\ \ r\ )^T$ denotes the circle center and radius. Consider a collection of points, $\boldsymbol{X_i}$, $i = 1, \ldots, m$, each of which satisfy Equation (3). Under an ideal pinhole camera model, the world point $\boldsymbol{X_i}$ maps to the image point $\boldsymbol{x_i}$ as:

$$\boldsymbol{x_i} = H\boldsymbol{X_i}, \tag{4}$$

where H is the 3×3 homography.

The estimation of H can be formulated in an orthogonal distance fitting framework. Let $E(\cdot)$ be an error function on the parameter vector $\boldsymbol{\alpha}$ and the unknown homography H:

$$E(\boldsymbol{\alpha}, H) = \sum_{i=1}^{m} \min_{\hat{\boldsymbol{X}}} \left\| \boldsymbol{x_i} - H\hat{\boldsymbol{X}} \right\|^2, \tag{5}$$

where $\hat{\boldsymbol{X}}$ is on the circle parametrized by $\boldsymbol{\alpha}$. This error embodies the sum of the squared errors between the data, $\boldsymbol{x_i}$, and the closest point on the model, $\hat{\boldsymbol{X}}$. One circle provides five constraints on the nine unknowns of H. In order to estimate H, at least one other circle is required. With two circles, the above error function takes the form:

$$\begin{aligned}\hat{E}(\boldsymbol{\alpha}_1, H_1, \boldsymbol{\alpha}_2, H_2) &= E(\boldsymbol{\alpha}_1, H_1) + E(\boldsymbol{\alpha}_2, H_2) \\ &\quad + \omega \left(\|\boldsymbol{H}_1 - \boldsymbol{H}_2\|^2 + (r_1 - 5.8)^2 + (r_2 - 5.8)^2 \right),\end{aligned} \tag{6}$$

where w is a scalar weighting factor. The first two terms are the individual error functions, Equation (5), for the two circles. The remaining terms constrain the transforms for both circles to be the same,[1] and the radius to be equal to 5.8 mm. This error function is minimized using non-linear least squares via the Levenberg-Marquardt iteration [6]. When the two circles are co-planar with respect to the camera plane, the eyes will be imaged as circles regardless of where they are in the world coordinate system. In this case, the principal point cannot be uniquely determined, and is assumed to be at the image center.

2.2 Camera Calibration

Once estimated, the homography H can be decomposed in terms of its intrinsic and extrinsic camera parameters [16, 18]. The intrinsic parameters consist of the focal length f, principal point (c_1, c_2), skew σ, and aspect ratio α. The extrinsic parameters consist of a rotation matrix R and translation vector $\boldsymbol{t}$ that define

[1] The notation $\boldsymbol{H}_i$ expresses the matrix H_i as a vector.

the transformation between the world and camera coordinate systems. Since the world points lie on a single plane, H can be decomposed in terms of the intrinsic and extrinsic parameters [3] as:

$$H = \lambda K \left(\boldsymbol{r}_1 \; \boldsymbol{r}_2 \; \boldsymbol{t} \right), \tag{7}$$

where λ is a scale factor and the 3×3 intrinsic matrix K is:

$$K = \begin{pmatrix} \alpha f & \sigma & c_1 \\ 0 & f & c_2 \\ 0 & 0 & 1 \end{pmatrix}. \tag{8}$$

For simplicity, we will assume that the skew (σ) is zero and that the aspect ratio (α) is 1. Under these assumptions, the matrix K is:

$$K = \begin{pmatrix} f & 0 & c_1 \\ 0 & f & c_2 \\ 0 & 0 & 1 \end{pmatrix}. \tag{9}$$

The camera's intrinsic components can be estimated by decomposing H according to Equation (7). It is straightforward to show that that $\boldsymbol{r}_1 = \frac{1}{\lambda} K^{-1} \boldsymbol{h}_1$ and $\boldsymbol{r}_2 = \frac{1}{\lambda} K^{-1} \boldsymbol{h}_2$ where $\boldsymbol{h}_1$ and $\boldsymbol{h}_2$ are the first two columns of the matrix H. The constraint that $\boldsymbol{r}_1$ and $\boldsymbol{r}_2$ are orthogonal (they are columns of a rotation matrix) and have the same norm (unknown due to the scale factor λ) yields two constraints on the unknown matrix K:

$$\boldsymbol{r}_1^T \boldsymbol{r}_2 = \boldsymbol{h}_1^T (K^{-T} K^{-1}) \boldsymbol{h}_2 = 0, \tag{10}$$

$$\boldsymbol{r}_1^T \boldsymbol{r}_1 - \boldsymbol{r}_2^T \boldsymbol{r}_2 = \boldsymbol{h}_1^T (K^{-T} K^{-1}) \boldsymbol{h}_1 - \boldsymbol{h}_2^T (K^{-T} K^{-1}) \boldsymbol{h}_2 = 0. \tag{11}$$

With only two constraints, it is possible to estimate the principal point (c_1, c_2) or the focal length f, but not both [18]. As such, we will assume a known focal length.

For notational simplicity we solve for the components of $Q = K^{-T} K^{-1}$, which contain the desired coordinates of the principal point and the assumed known focal length:

$$Q = \frac{1}{f^2} \begin{pmatrix} 1 & 0 & -c_1 \\ 0 & 1 & -c_2 \\ -c_1 & -c_2 & c_1^2 + c_2^2 + f^2 \end{pmatrix}. \tag{12}$$

In terms of Q, the first constraint, Equation (10), takes the form:

$$\begin{aligned} h_1 h_2 + h_4 h_5 - (h_2 h_7 + h_1 h_8) c_1 - (h_5 h_7 + h_4 h_8) c_2 \\ + h_7 h_8 (c_1^2 + c_2^2 + f^2) = 0, \end{aligned} \tag{13}$$

where h_i is the i^{th} element of the matrix H in row-major order. Note that this constraint is a second-order polynomial in the coordinates of the principal point, which can be factored as follows:

$$(c_1 - \alpha_1)^2 + (c_2 - \beta_1)^2 = \gamma_1^2, \tag{14}$$

where:

$$\alpha_1 = (h_2 h_7 + h_1 h_8)/(2 h_7 h_8), \tag{15}$$
$$\beta_1 = (h_5 h_7 + h_4 h_8)/(2 h_7 h_8), \tag{16}$$
$$\gamma_1^2 = \alpha_1^2 + \beta_1^2 - f^2 - (h_1 h_2 + h_4 h_5)/(h_7 h_8). \tag{17}$$

Similarly, the second constraint, Equation (11), takes the form:

$$h_1^2 + h_4^2 + 2(h_2 h_8 - h_1 h_7)c_1 + 2(h_5 h_8 - h_4 h_7)c_2 - h_2^2 - h_5^2 + (h_7^2 - h_8^2)(c_1^2 + c_2^2 + f^2) = 0, \tag{18}$$

or,

$$(c_1 - \alpha_2)^2 + (c_2 - \beta_2)^2 = \gamma_2^2, \tag{19}$$

where:

$$\alpha_2 = (h_1 h_7 - h_2 h_8)/(h_7^2 - h_8^2), \tag{20}$$
$$\beta_2 = (h_4 h_7 - h_5 h_8)/(h_7^2 - h_8^2), \tag{21}$$
$$\gamma_2^2 = \alpha_2^2 + \beta_2^2 - (h_1^2 + h_4^2 - h_2^2 - h_5^2)/(h_7^2 - h_8^2) - f^2. \tag{22}$$

Both constraints, Equations (14) and (19) are circles in the desired coordinates of the principal point c_1 and c_2, and the solution is the intersection of the two circles.[2]

For certain homographies, however, this solution can be numerically unstable. For example, if $h_7 \approx 0$ or $h_8 \approx 0$, the first constraint becomes numerically unstable. Similarly, if $h_7 \approx h_8$, the second constraint becomes unstable. In order to avoid these instabilities, an error function with a regularization term is introduced.

We start with the following error function to be minimized:

$$E(c_1, c_2) = g_1(c_1, c_2)^2 + g_2(c_1, c_2)^2, \tag{23}$$

where $g_1(c_1, c_2)$ and $g_2(c_1, c_2)$ are the constraints on the principal point given in Equations (13) and (18), respectively. To avoid numerical instabilities, a regularization term is added to penalize deviations of the principal point from the image center $(0, 0)$ (in normalized coordinates). This augmented error function takes the form:

$$E(c_1, c_2) = g_1(c_1, c_2)^2 + g_2(c_1, c_2)^2 + \Delta(c_1^2 + c_2^2), \tag{24}$$

where Δ is a scalar weighting factor. This error function is a nonlinear least-squares problem, which can be minimized using a Levenberg-Marquardt iteration. The image center $(0, 0)$ is used as the initial condition for the iteration.

[2] In fact, there can be zero, one, two or an infinite number of real solutions depending on the configuration of the circles.

2.3 Translation

The translation of two circles (eyes) in the image is equivalent to translating the camera's principal point. In homogeneous coordinates, translations are represented by multiplication with a translation matrix T:

$$\boldsymbol{y} = T\boldsymbol{x}, \tag{25}$$

where:

$$T = \begin{pmatrix} 1 & 0 & d_1 \\ 0 & 1 & d_2 \\ 0 & 0 & 1 \end{pmatrix}, \tag{26}$$

and the amount of translation is (d_1, d_2). The mapping from world $\boldsymbol{X}$ to (translated) image coordinates $\boldsymbol{y}$ is:

$$\begin{aligned} \boldsymbol{y} &= TH\boldsymbol{X} \\ &= \lambda TK \begin{pmatrix} \boldsymbol{r}_1 & \boldsymbol{r}_2 & \boldsymbol{t} \end{pmatrix} \boldsymbol{X} \\ &= \lambda \hat{K} \begin{pmatrix} \boldsymbol{r}_1 & \boldsymbol{r}_2 & \boldsymbol{t} \end{pmatrix} \boldsymbol{X}, \end{aligned} \tag{27}$$

where

$$\hat{K} = \begin{pmatrix} f & 0 & c_1 + d_1 \\ 0 & f & c_2 + d_2 \\ 0 & 0 & 1 \end{pmatrix}. \tag{28}$$

Therefore, translation in image coordinates is equivalent to translating the principal point. Assuming the principal point in an authentic image is near the origin [17], large deviations from the image center, or inconsistencies in the estimated principal point across the image, can be used as evidence of tampering.

3 Results

We tested our technique for estimating the principal point from images of eyes on synthetically generated images, real images, and visually plausible forgeries. In all of these results, the principal point was estimated by minimizing Equation (24). Because of the regularization term used in this solution, we found that the estimated principal point was biased towards the image center $(0, 0)$. For purely aesthetic purposes, we rescaled the norm, n, of the estimated principal point by $3n^{1.7}$, where the form of this correction was chosen empirically. Throughout, we will refer to normalized image coordinates where the image center is $(0, 0)$ and the horizontal and vertical coordinates are normalized so that the maximum of the dimensions is in the range $[-1, 1]$.

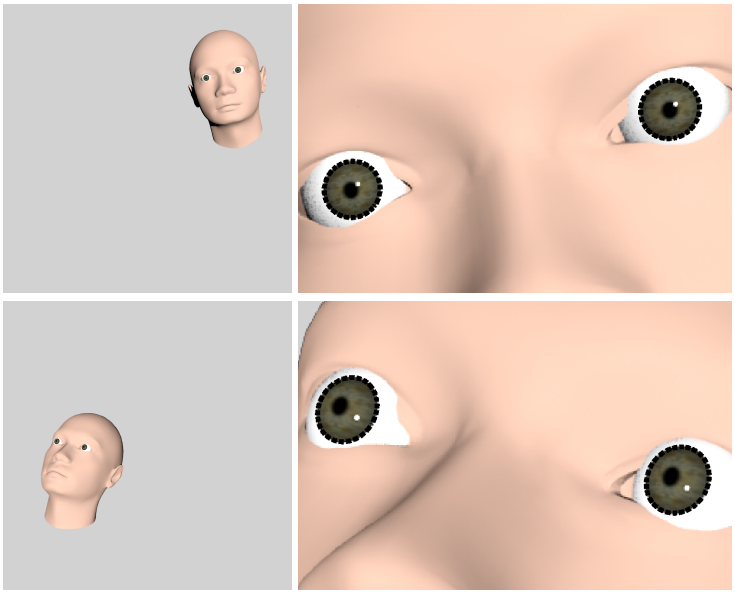

Fig. 3. A 3-D model of a head in two different locations and orientations (left), and a magnified view of the eyes with the extracted boundaries of the limbi (right).

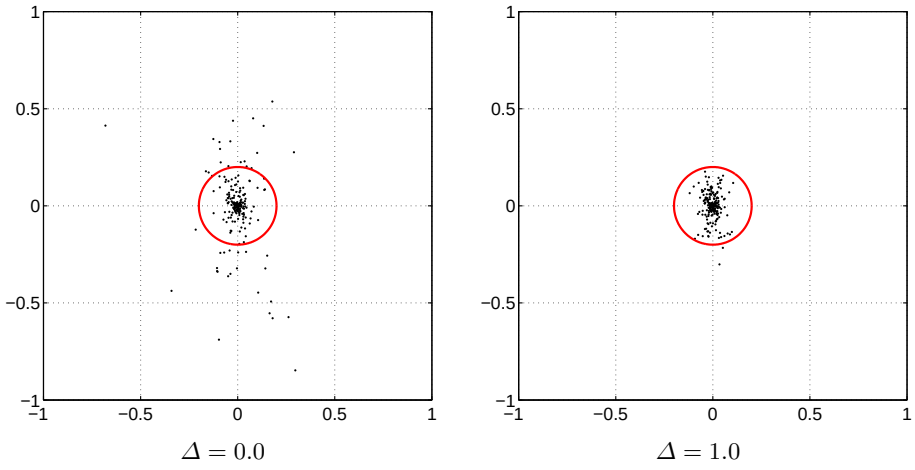

Fig. 4. Estimates of the principal point in normalized coordinates for 216 synthetically generated images, Fig 3. Shown on the left are the unconditioned estimates ($\Delta = 0.0$) and on the right are the conditioned estimates ($\Delta = 1.0$). Note that the conditioning significantly improves the accuracy of the estimation – the actual principal point is the origin $(0, 0)$. A circle at a threshold of 0.2 units is drawn for reference.

3.1 Synthetic

Shown in Fig. 3 are two examples of synthetically generated heads rendered using the `pbrt` environment [12]. Also shown in this figure is a magnified view of the eyes whose shape conformed to the model described in Section 2.1. The eyes were rendered with a full head model to provide a notion of size, though only the shape of the limbi are used in the subsequent analysis. Each image was rendered at 2400×2400 pixels and the radius of the limbus ranged from 24 to 34 pixels.

In the first set of simulations, the head model was rotated to 8 different orientations (ranging from -15 to 15 degrees from parallel) and 27 different locations, for a total of 216 images. Two sample images are shown in Fig. 3. The elliptical shapes of the limbi in each image were automatically extracted, Fig. 3. The homography H was estimated as described in the previous section, with a regularization term $\Delta = 0.0$ or $\Delta = 1.0$, Equation (24). The actual principal point in these synthetically generated images is the image center: $(0, 0)$ in normalized coordinates.

Shown in the left panel of Fig. 4 are the unconditioned estimates ($\Delta = 0.0$) and in the right panel are the conditioned estimates ($\Delta = 1.0$). Note that the conditioning significantly improves the accuracy of the estimation: without conditioning 80.6% (174/216) of the estimates are within 0.2 units of the origin (red circle), and with conditioning 99.1% (214/216) are within 0.2 units of the origin.

In the second set of simulations, the head model was positioned at the center of the world coordinate system and rotated to 252 different orientations. The rotations ranged from -30 to 30 degrees about each axis. Shown in Fig. 5 are four sample images. To simulate tampering, the head in each image was translated to various locations and the principal point was estimated at each location. Superimposed on the images in Fig. 5 are level curves denoting the deviation of the estimated principal point from the origin as a function of spatial position in the image. The inner curve denotes a distance of 0.2, and each subsequent curve denotes an increment of 0.1. With a threshold of 0.2, translations of the head outside of the inner-most curve can be detected as fake. Note that the level curves are typically asymmetric and depend on the orientation of the head.

In the third set of simulations, the 252 images from the previous experiment were translated in the image by random amounts such that the displacement was greater than 0.2 units in normalized coordinates (240 pixels in the original 2400×2400 image). Shown in the left panel of Fig. 6 are the estimated principal points for the original 252 images, and in the right panel are the estimated results for 1260 translated images. In both cases, the conditioned estimator ($\Delta = 1.0$) was used. Of the 252 authentic images, 99.2% had an estimated principal point less than 0.2 units from the origin, and of the 1260 translated images, 94.6% had an estimated principal point greater than 0.2 units from the origin.

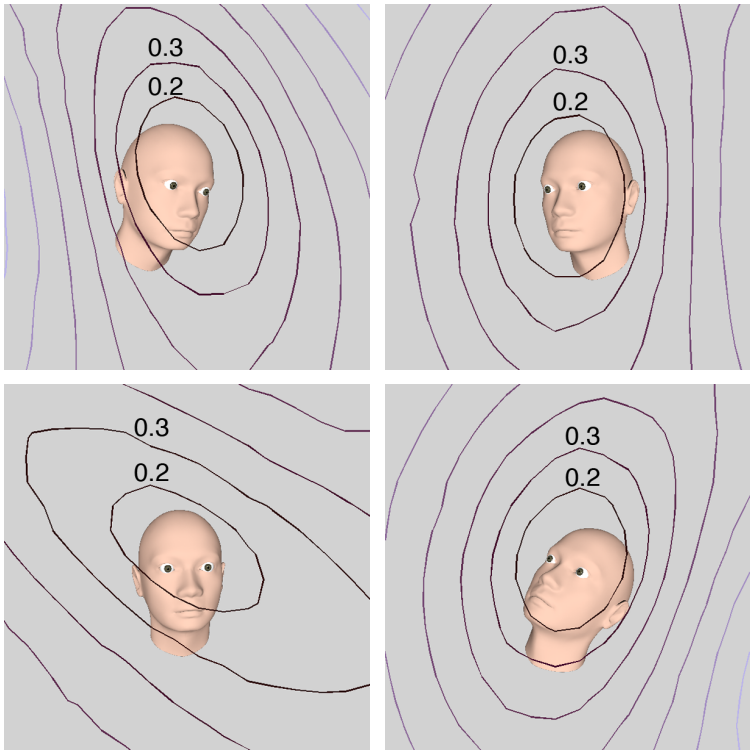

Fig. 5. A 3-D model of a head at four different orientations. The superimposed level curves show the deviation of the estimated principal point from the origin as a function of spatial position in the image.

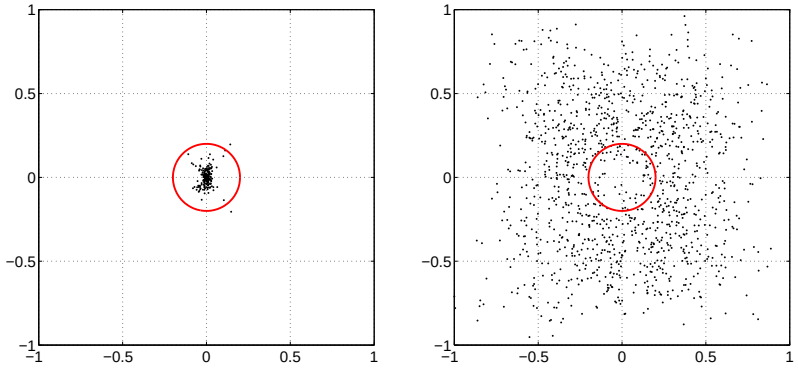

Fig. 6. Estimates of the principal point in normalized coordinates for 252 authentic images (left) and 1260 doctored images (right), Fig. 5. The actual principal point is the origin $(0, 0)$, and a circle at a threshold of 0.2 units is drawn for reference.

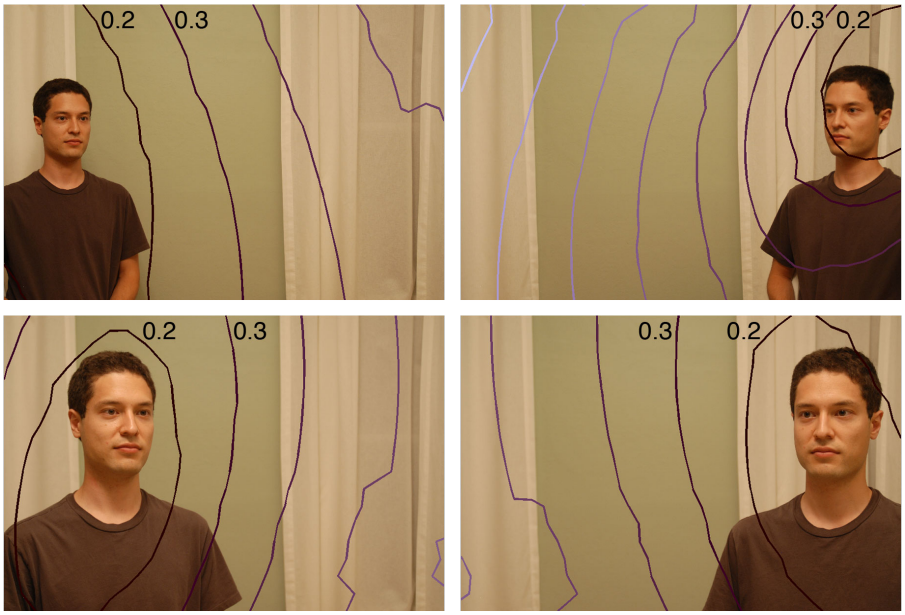

Fig. 7. Shown in each panel are authentic images with superimposed level curves showing the deviation of the estimated principal point from the true principal point as a function of spatial position in the image.

3.2 Real

Shown in Fig. 7 are four of 15 images taken with a Nikon D200 10-megapixel camera set to record at the highest quality JPEG format. At a radius of 21 pixels, the size of the eyes in these images was slightly smaller than in the above simulations. The principal point was estimated using the conditioned estimator, Equation (24), with $\Delta = 1.0$. The average deviation from the calibrated principal point[3] was 0.15 units (in normalized coordinates) with a maximum distance of 0.24, a minimum distance of 0.05 and a standard deviation of 0.06 units. The eyes in each of the four images were translated to various locations in the image to simulate tampering. The level curves in Fig. 7 show the deviation of the estimated principal point from the true principal point, as a function of spatial position in the image. With a threshold of 0.2 units, translations in the image outside of the innermost curve are classified as fake.

[3] The camera was calibrated to determine the actual principal point which at $(-0.028, 0.022)$ is close to the origin. The Camera Calibration Toolbox for Matlab `http://www.vision.caltech.edu/bouguetj/calib_doc` was used for this calibration.

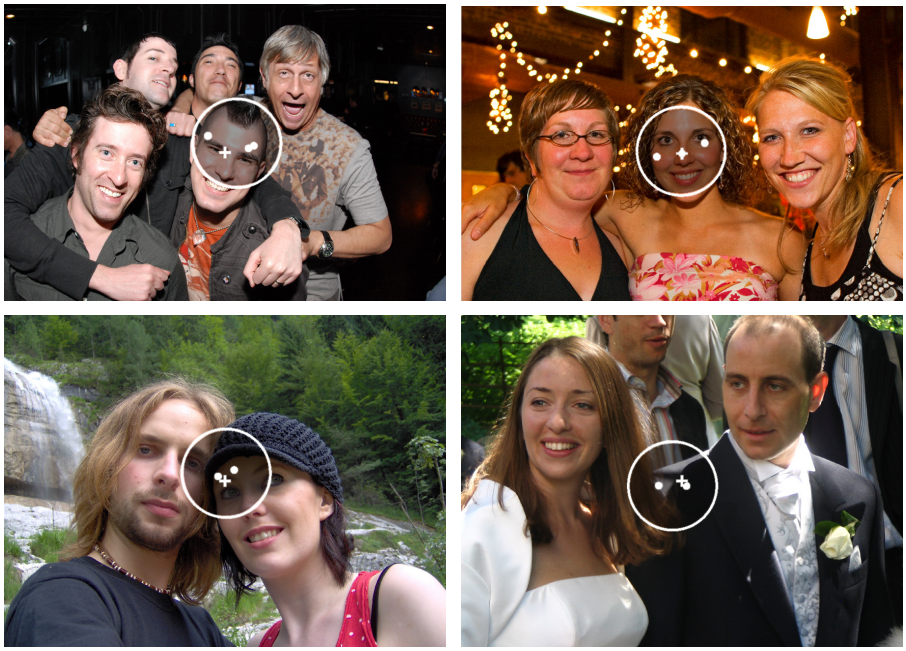

Fig. 8. Four authentic images. In each image, the '+' marker denotes the image center $(0, 0)$, and the dots denote the estimated principal points from each person (the eyes from only three people in the top left image were visible). The circle with radius 0.2 units is centered at the average of the principal points.

Shown in Fig. 8 are four images acquired from Flickr, a popular image sharing website. The images were captured with different cameras at different focal lengths. The focal length was extracted from the metadata in the image and used to estimate the principal point. In each image, the '+' marker denotes the image center $(0, 0)$ and the white dots denote the principal points estimated from different people. The circle in each image has a radius of 0.2 and is centered at the average of the principal points. Note that in each of the four images, the estimated principal points fall within the circle, indicating relative agreement in the positions of the camera's principal point.

Since the homography can be estimated from other known geometries [5], the estimation of the principal point is not limited to images of the eyes. Shown in Fig. 9, for example, are results from a car's wheels, and the known geometry of a stop sign.

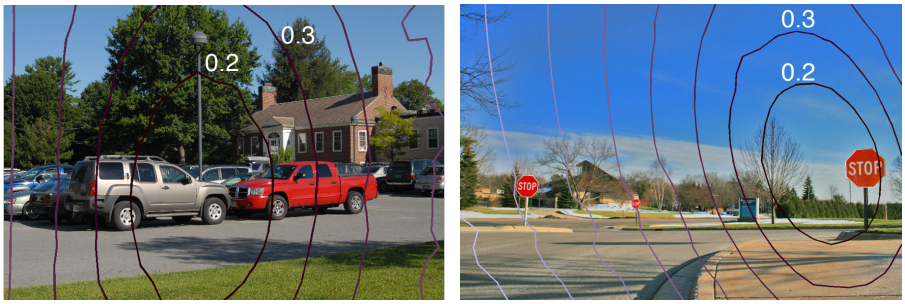

Fig. 9. Shown in each panel are authentic images with superimposed level curves showing the deviation of the estimated principal point from the true principal point as a function of spatial position in the image. The principal point for the left image was determined from the left-most car's wheels, and for the right image, the known geometry of the stop sign.

3.3 Forgeries

We created two forgeries by combining images from Fig. 7 and Fig. 8. As shown in Fig. 10, the principal points estimated from the forged heads are inconsistent with the other principal point(s) in the image.

4 Discussion

When creating a composite of two or more people it is often necessary to move a person in the image relative to their position in the original image. When done well, this manipulation is rarely visually obvious. We have shown how to detect such manipulations by estimating the camera's principal point (the projection of the camera center onto the image plane) from the image of a person's eyes. This approach relies on estimating the transformation from world to image coordinates and then factoring this transformation into a product of matrices containing intrinsic and extrinsic camera parameters. With a known focal length, the principal point can be determined from the intrinsic matrix. Inconsistencies in the estimated principal point can then be used as evidence of tampering.

We have shown the efficacy of this technique on simulated and real images. The major sensitivity with this technique is in extracting the elliptical boundary of the eye. This process will be particularly difficult for low-resolution images, but with a radius of $20 - 30$ pixels reasonably accurate estimates can be made from a person's eyes.

We expect this technique, in conjunction with a growing body of forensic tools, to be effective in exposing digital forgeries.

Fig. 10. Two forgeries made by combining images from Fig. 7 and Fig. 8. The '+' marker denotes the image center $(0, 0)$, and the dots denote the estimated principal points from each person. The circle with radius 0.2 units is centered at the average of the principal points. Notice that the estimated principal points are inconsistent with one another.

Acknowledgments

This work was supported by a Guggenheim Fellowship, a gift from Adobe Systems, Inc., a gift from Microsoft, Inc., a grant from the United States Air Force (FA8750-06-C-0011), and by the Institute for Security Technology Studies at Dartmouth College under grant 2005-DD-BX-1091 from the Bureau of Justice Assistance and Award Number 2006-CS-001-000001 from the U.S. Department of Homeland Security. Points of view or opinions in this document are those of the author and do not represent the official position or policies of the U.S. Department of Justice, the U.S. Department of Homeland Security, or any other sponsor.

References

1. İsmail Avcıbaş, Sevinç Bayram, Nasir Memon, Bülent Sankur, and Mahalingam Ramkumar. A classifier design for detecting image manipulations. In *2004 International Conference on Image Processing, ICIP '04*, volume 4, pages 2645–2648, 2004.
2. Jessica Fridrich, David Soukal, and Jan Lukáš. Detection of copy-move forgery in digital images. In *Proceedings of Digital Forensic Research Workshop*, August 2003.
3. Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.
4. Micah K. Johnson and Hany Farid. Exposing digital forgeries by detecting inconsistencies in lighting. In *ACM Multimedia and Security Workshop*, 2005.
5. Micah K. Johnson and Hany Farid. Metric measurements on a plane from a single image. Technical Report TR2006-579, Department of Computer Science, Dartmouth College, 2006.

6. Micah K. Johnson and Hany Farid. Exposing digital forgeries through specular highlights on the eye. In *9th International Workshop on Information Hiding*, Saint Malo, France, 2007.
7. Micah K. Johnson and Hany Farid. Exposing digital forgeries in complex lighting environments. *IEEE Transactions on Information Forensics and Security*, 2007 (in press).
8. Aaron Lefohn, Richard Caruso, Erik Reinhard, Brian Budge, and Peter Shirley. An ocularist's approach to human iris synthesis. *IEEE Computer Graphics and Applications*, 23(6):70–75, 2003.
9. David Liebowitz and Andrew Zisserman. Metric rectification for perspective images of planes. In *Computer Vision and Pattern Recognition*, pages 482–488, 1998.
10. Jan Lukáš, Jessica Fridrich, and Miroslav Goljan. Detecting digital image forgeries using sensor pattern noise. In *Proceedings of the SPIE*, volume 6072, 2006.
11. Tian-Tsong Ng and Shih-Fu Chang. A model for image splicing. In *IEEE International Conference on Image Processing (ICIP)*, Singapore, October 2004.
12. Matt Pharr and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann, 2004.
13. Alin C. Popescu and Hany Farid. Exposing digital forgeries by detecting duplicated image regions. Technical Report TR2004-515, Department of Computer Science, Dartmouth College, 2004.
14. Alin C. Popescu and Hany Farid. Exposing digital forgeries by detecting traces of re-sampling. *IEEE Transactions on Signal Processing*, 53(2):758–767, 2005.
15. Alin C. Popescu and Hany Farid. Exposing digital forgeries in color filter array interpolated images. *IEEE Transactions on Signal Processing*, 53(10):3948–3959, 2005.
16. Roger Y. Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf cameras and lenses. *IEEE Journal of Robotics and Automation*, RA-3(4):323–344, August 1987.
17. Reg G. Willson and Steven A. Shafer. What is the center of the image? *Journal of the Optical Society of America A*, 11(11):2946–2955, November 1994.
18. Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1330–1334, 2000.