

Exposing Digital Forgeries by Detecting Inconsistencies in Lighting

Micah K. Johnson
Department of Computer Science
Dartmouth College
Hanover, NH 03755
kimo@cs.dartmouth.edu

Hany Farid
Department of Computer Science
Dartmouth College
Hanover, NH 03755
farid@cs.dartmouth.edu

ABSTRACT

When creating a digital composite of, for example, two people standing side-by-side, it is often difficult to match the lighting conditions from the individual photographs. Lighting inconsistencies can therefore be a useful tool for revealing traces of digital tampering. Borrowing and extending tools from the field of computer vision, we describe how the direction of a point light source can be estimated from only a single image. We show the efficacy of this approach in real-world settings.

Categories and Subject Descriptors

I.4 [Image Processing]: Miscellaneous

Keywords

Digital Tampering, Digital Forensics

1. INTRODUCTION

Consider the creation of a forgery showing two movie stars, rumored to be romantically involved, walking down a sunset beach. Such an image might be created by splicing together individual images of each movie star. In so doing, it is often difficult to exactly match the lighting effects due to directional lighting (e.g., the sun on a clear day). Differences in lighting can, therefore, be a telltale sign of digital tampering. Shown in Figure 1, for example, is a composite image where the two people were originally photographed with the light in significantly different positions. While this type of forgery is fairly obvious, more subtle differences in lighting direction may be harder to detect by simple visual inspection [13, 7].

To the extent that the direction of the light source can be estimated for different objects/people in an image, inconsistencies in the lighting direction can be used as evidence of digital tampering. In this paper, we describe a technique



Figure 1: A digital composite of movie stars Cher and Brad Pitt. Note that Cher was originally photographed with a fairly diffuse non-directional light source, whereas Brad Pitt was photographed with a directional light positioned to his left.

for estimating the light source direction from a single image, and show its efficacy in real-world settings.

2. METHODS

The general problem of estimating the illuminant direction has been widely studied in the field of computer vision (e.g., [8, 2, 6]). In this section, we define the general problem, review a standard solution and then show how some additional simplifying assumptions make the problem more tractable. We then extend this solution to provide for a more effective and broadly applicable forensic tool.

2.1 Infinite Light Source (3-D)

The standard approaches for estimating light source direction begin by making some simplifying assumptions: (1) the surface of interest is Lambertian (the surface reflects light isotropically); (2) the surface has a constant reflectance value; (3) the surface is illuminated by a point light source infinitely far away; and (4) the angle between the surface normal and the light direction¹ is in the range 0° to 90° .

¹The assumption that the angle between the surface and light is bounded between 0° to 90° can be relaxed by replacing $(\vec{N}(x, y) \cdot \vec{L})$ in Equation (1) with $\max(\vec{N}(x, y) \cdot \vec{L}, 0)$, which is not used here to avoid the non-linear max operator.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM Multimedia and Security Workshop '05 New York, New York USA
Copyright 2005 ACM 1-59593-032-9/05/0008 ...\$5.00.

Under these assumptions, the image intensity can be expressed as:

$$I(x, y) = R(\vec{N}(x, y) \cdot \vec{L}) + A, \quad (1)$$

where R is the constant reflectance value, \vec{L} is a 3-vector pointing in the direction of the light source, $\vec{N}(x, y)$ is a 3-vector representing the surface normal at the point (x, y) , and A is a constant ambient light term [3], Figure 2(a). If we are only interested in the direction of the light source, then the reflectance term, R , can be considered to have unit-value, understanding that the estimation of \vec{L} will only be within an unknown scale factor. The resulting linear equation provides a single constraint in four unknowns, the three components of \vec{L} and the ambient term A .

With at least four points with the same reflectance, R , and distinct surface normals, \vec{N} , the light source direction and ambient term can be solved for using standard least-squares estimation. To begin, a quadratic error function, embodying the imaging model of Equation (1), is given by:

$$\begin{aligned} E(\vec{L}, A) &= \left\| M \begin{pmatrix} L_x \\ L_y \\ L_z \\ A \end{pmatrix} - \begin{pmatrix} I(x_1, y_1) \\ I(x_2, y_2) \\ \vdots \\ I(x_p, y_p) \end{pmatrix} \right\|^2 \\ &= \left\| M\vec{v} - \vec{b} \right\|^2, \end{aligned} \quad (2)$$

where $\|\cdot\|$ denotes vector norm, L_x , L_y , and L_z denote the components of the light source direction \vec{L} , and

$$M = \begin{pmatrix} N_x(x_1, y_1) & N_y(x_1, y_1) & N_z(x_1, y_1) & 1 \\ N_x(x_2, y_2) & N_y(x_2, y_2) & N_z(x_2, y_2) & 1 \\ \vdots & \vdots & \vdots & \vdots \\ N_x(x_p, y_p) & N_y(x_p, y_p) & N_z(x_p, y_p) & 1 \end{pmatrix}, \quad (3)$$

where $N_x(x_i, y_i)$, $N_y(x_i, y_i)$, and $N_z(x_i, y_i)$ denote the components of the surface normal \vec{N} at image coordinate (x_i, y_i) . The quadratic error function above is minimized by differentiating with respect to the unknown, \vec{v} , setting the result equal to zero, and solving for \vec{v} to yield the least-squares estimate:

$$\vec{v} = (M^T M)^{-1} M^T \vec{b}. \quad (4)$$

Note that this solution requires knowledge of 3-D surface normals from at least four distinct points ($p \geq 4$) on a surface with the same reflectance. With only a single image and no objects of known geometry in the scene, it is unlikely that this will be possible. Most approaches to overcome this problem rely on acquiring multiple images [9] or placing an object of known geometry in the scene (e.g., a sphere) [1]. For forensic applications, these solutions are not practical.

2.2 Infinite Light Source (2-D)

In [6], the authors suggest a clever solution for estimating two components of the light source direction (L_x and L_y) from only a single image. While their approach clearly provides less information regarding the light source direction, it does make the problem tractable from a single image. The authors note that at the occluding boundary of a surface, the z -component of the surface normal is zero, $N_z = 0$. In addition, the x - and y -components of the surface normal, N_x and N_y , can be estimated directly from the image, Figure 2(b).

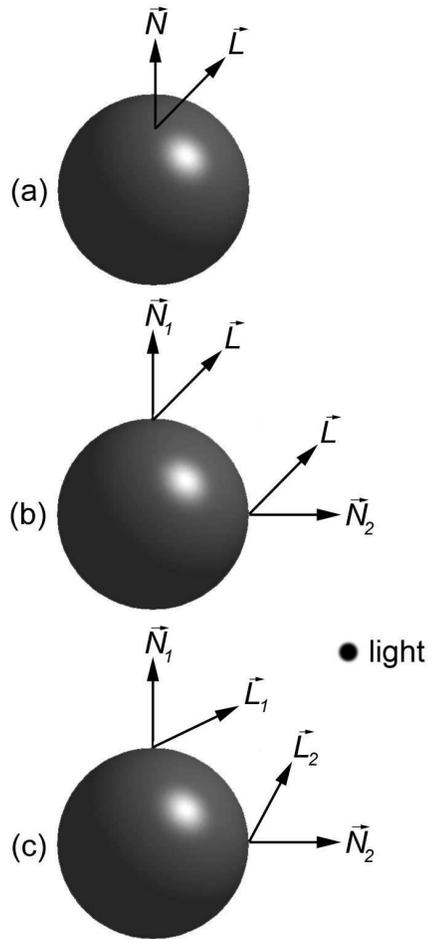


Figure 2: Schematic diagram of the imaging geometry for an (a) infinite light source (3-D); (b) infinite light source (2-D); and (c) local light source (2-D). In the 2-D cases, the z -component of the surface normal (\vec{N}) is zero. Unlike an infinite light source, the direction to a local light source (\vec{L}) varies across the sphere's surface.

With this assumption, the error function of Equation (2) takes the form:

$$\begin{aligned} E(\vec{L}, A) &= \left\| M \begin{pmatrix} L_x \\ L_y \\ A \end{pmatrix} - \begin{pmatrix} I(x_1, y_1) \\ I(x_2, y_2) \\ \vdots \\ I(x_p, y_p) \end{pmatrix} \right\|^2 \\ &= \left\| M\vec{v} - \vec{b} \right\|^2, \end{aligned} \quad (5)$$

where,

$$M = \begin{pmatrix} N_x(x_1, y_1) & N_y(x_1, y_1) & 1 \\ N_x(x_2, y_2) & N_y(x_2, y_2) & 1 \\ \vdots & \vdots & \vdots \\ N_x(x_p, y_p) & N_y(x_p, y_p) & 1 \end{pmatrix}. \quad (6)$$

This error function is minimized, as before, using standard least-squares to yield the same solution as in Equation (4),

but with the matrix M taking the form given in Equation (6). In this case, the solution requires knowledge of 2-D surface normals from at least three distinct points ($p \geq 3$) on a surface with the same reflectance.

The intensity, $I(x_i, y_i)$, at a boundary point, (x_i, y_i) , cannot be directly measured from the image as the surface is occluded. The authors in [6] note, however, that the intensity can be extrapolated by considering the intensity profile along a ray coincident to the 2-D surface normal. They also found that simply using the intensity close to the border of the surface is often sufficient (see Section 3 for a more detailed description).

We extend this basic formulation in three ways. First, we estimate the two-dimensional light source direction from local patches along an object's boundary (as opposed to along extended boundaries as in [6]). This is done to relax the assumption that the reflectance along the entire surface is constant, Section 2.2.1. Then, we introduce a regularization (smoothness) term to better condition the final estimate of light source direction. Finally, this formulation is extended to accommodate a local directional light source (e.g., a desk lamp), Section 2.3.

2.2.1 Relaxing the Constant Reflectance Assumption

We relax the constant reflectance assumption by assuming that the reflectance for a local surface patch (as opposed to the entire surface) is constant. This requires us to estimate individual light source directions, \vec{L}^i , for each patch along a surface. Under the infinite light source assumption, the orientation of these estimates should not vary, but their magnitude may (recall that the estimate of the light source is only within a scale factor, which depends on the reflectance value R , Equation (1)).

Consider a surface partitioned into n patches, and, for notational simplicity, assume that each patch contains p points. The new error function to be minimized is constructed by packing together, for each patch, the 2-D version of the constraint of Equation (1):

$$E_1(\vec{L}^1, \dots, \vec{L}^n, A) = \left\| M \begin{pmatrix} L_x^1 \\ L_y^1 \\ \vdots \\ L_x^n \\ L_y^n \\ A \end{pmatrix} - \begin{pmatrix} I(x_1^1, y_1^1) \\ \vdots \\ I(x_p^1, y_p^1) \\ \vdots \\ I(x_1^n, y_1^n) \\ \vdots \\ I(x_p^n, y_p^n) \end{pmatrix} \right\|^2$$

$$= \left\| M\vec{v} - \vec{b} \right\|^2, \quad (7)$$

where,

$$M = \begin{pmatrix} N_x(x_1^1, y_1^1) & N_y(x_1^1, y_1^1) & 0 & 0 & 1 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ N_x(x_p^1, y_p^1) & N_y(x_p^1, y_p^1) & 0 & 0 & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & N_x(x_1^n, y_1^n) & N_y(x_1^n, y_1^n) & 1 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & N_x(x_p^n, y_p^n) & N_y(x_p^n, y_p^n) & 1 \end{pmatrix} \quad (8)$$

The above quadratic error function is minimized, as before, using least-squares with the solution taking on the same

form as in Equation (4). In this case, the solution provides n estimates of the 2-D light directions, $\vec{L}^1, \dots, \vec{L}^n$, and an ambient term A . Note that while individual light source directions are estimated for each surface patch, a single ambient term is assumed.

While the local estimation of light source directions allows for the relaxation of the constant reflectance assumption, it could potentially yield less stable results. Note that under the assumption of an infinite point light source, the orientation of the n light directions should be equal. With the additional assumption that the change in reflectance from patch to patch is relatively small (i.e., the change in the magnitude of neighboring \vec{L}^i 's is small), we can condition the individual estimates with the following regularization term:

$$E_2(\vec{L}^1, \dots, \vec{L}^n) = \sum_{i=2}^n \left\| \vec{L}^i - \vec{L}^{i-1} \right\|^2. \quad (9)$$

This additional error term penalizes neighboring estimates that are different from one another. The quadratic error function $E_1(\cdot)$, Equation (7), is conditioned by combining it with the regularization term $E_2(\cdot)$, scaled by a factor λ , to yield the final error function:

$$E(\vec{L}^1, \dots, \vec{L}^n, A) = E_1(\vec{L}^1, \dots, \vec{L}^n, A) + \lambda E_2(\vec{L}^1, \dots, \vec{L}^n). \quad (10)$$

This combined error function can still be minimized using least-squares minimization. The error function $E_2(\cdot)$ is first written in a more compact and convenient form as:

$$E_2(\vec{v}) = \|C\vec{v}\|^2, \quad (11)$$

where the $2n - 2 \times 2n + 1$ matrix C is given by:

$$C = \begin{pmatrix} -1 & 0 & 1 & 0 & \dots & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 & \dots & 0 & 0 & 0 & 0 & 0 \\ \vdots & & & & \ddots & & & & & \vdots \\ 0 & 0 & 0 & 0 & \dots & -1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & -1 & 0 & 1 & 0 \end{pmatrix}, \quad (12)$$

with $\vec{v} = (L_x^1 \ L_y^1 \ L_x^2 \ L_y^2 \ \dots \ L_x^n \ L_y^n \ A)^T$. The error function of Equation (10) then takes the form:

$$E(\vec{v}) = \|M\vec{v} - \vec{b}\|^2 + \lambda \|C\vec{v}\|^2. \quad (13)$$

Differentiating this error function yields:

$$E'(\vec{v}) = 2M^T M\vec{v} - 2M^T \vec{b} + 2\lambda C^T C\vec{v}$$

$$= 2(M^T M + \lambda C^T C)\vec{v} - 2M^T \vec{b}. \quad (14)$$

Setting this result equal to zero and solving for \vec{v} yields the least-squares estimate:

$$\vec{v} = (M^T M + \lambda C^T C)^+ M^T \vec{b}, \quad (15)$$

where $+$ denotes pseudo-inverse. The final light direction estimate is computed by averaging the n resulting light direction estimates, $\vec{L}^1, \dots, \vec{L}^n$.

2.3 Local Light Source (2-D)

Inherent to the formulation of the previous two sections was the assumption that the directional light source, \vec{L} , was infinitely far away (i.e., \vec{L} does not depend on the image coordinates). With a local light source, however, this assumption is no longer valid, Figure 2(c). The model for

an infinite light source, Equation (1), can be rewritten to accommodate a local light source as follows:

$$I(x, y) = R(\vec{N}(x, y) \cdot \vec{L}(x, y)) + A. \quad (16)$$

Note that the light source direction is now a function of the image coordinates.

We begin by assuming that the light source direction for a local surface patch is constant across the patch. The light source direction for each surface patch is then estimated using the solution of Equation (7). In the previous section, a regularization term, that encouraged neighboring estimates to be equal, was introduced, Equation (9). In the case of a local light source, a different regularization term is needed as neighboring directions are expected to converge to a single nearby point. This regularization term takes the form:

$$E_2(\vec{L}^1, \dots, \vec{L}^n) = \sum_{i=1}^n \left\| C_i \vec{L}^i \right\|^2, \quad (17)$$

where the matrix C_i is derived in Appendix A. As in the previous section, the final error function to be minimized is given by:

$$E(\vec{L}^1, \dots, \vec{L}^n, A) = E_1(\vec{L}^1, \dots, \vec{L}^n, A) + \lambda E_2(\vec{L}^1, \dots, \vec{L}^n), \quad (18)$$

where $E_1(\cdot)$ is given by Equation (7), and λ is a scaling factor. Unlike the previous section, this error function cannot be minimized analytically, and is instead minimized using an iterative conjugate gradient minimization, Appendix B. Although the functional form of the error function appears similar to that of the previous section, the matrices C_i depend on the light source estimate \vec{L}^i , hence the need for an iterative minimization.

2.4 Multiple Light Sources

In the previous sections, it was assumed that a single directional light source was illuminating the scene (plus a constant ambient term). This is a reasonable assumption for outdoor images where the sun is typically the single source of illumination. For indoor images, however, this assumption is less reasonable because multiple light sources may be present.

Light has the wonderful property that it is linear. As such, a scene illuminated with, for example, two infinite light sources takes the form:

$$\begin{aligned} I(x, y) &= R((\vec{N}(x, y) \cdot \vec{L}_1) + (\vec{N}(x, y) \cdot \vec{L}_2)) + A \\ &= R(\vec{N}(x, y) \cdot (\vec{L}_1 + \vec{L}_2)) + A \\ &= R(\vec{N}(x, y) \cdot \vec{L}^+) + A, \end{aligned} \quad (19)$$

where \vec{L}^+ is the vector sum of the individual vectors \vec{L}_1 and \vec{L}_2 . Note that this model reduces to the same form as a single light source, Equation (1). Using the same approach as in the previous sections, therefore, will result in an estimate of a “virtual” light source, the vector sum of the individual light sources. This trivially extends to three or more individual light sources.

Although not terribly likely, it is possible that different combinations of light sources will sum to the same “virtual” light source, in which case this approach would be unable to detect an inconsistency in the lighting.

3. RESULTS

We tested our technique on both synthetically generated images and natural photographs. The synthetic images consisted of one or more spheres of constant reflectance rendered under either the infinite or local imaging models of Equation (1) or (16). The natural photographs were taken outdoors on a clear sunny day (approximating an infinite point light source), or in a controlled lab setting with a single directional light source (approximating a local point light source). These images were taken with a Nikon D-100 digital camera set to capture in uncompressed RAW format.

The light direction estimation requires the localization of an occluding boundary. These boundaries are extracted by manually selecting points in the image along an occluding boundary. This rough estimate of the position of the boundary is used to define its spatial extent. The boundary is then partitioned into approximately eight small patches. Three points near the occluding boundary are manually selected for each patch, and fit with a quadratic curve. The surface normals along each patch are then estimated analytically from the resulting quadratic fit.

The intensity from the occluding boundary cannot be directly measured from the image as the surface is occluded. The authors in [6] note, however, that simply using the intensity close to the border is often sufficient. The authors also found that under certain conditions it was advantageous to extrapolate the intensity by considering the intensity profile along a ray coincident to the 2-D surface normal. In the former case, the intensity is measured at a fixed number of pixels from the boundary in the direction opposite to the surface normal. More specifically, the intensity at a boundary point (x_i, y_i) with surface normal \vec{N} is determined by evaluating the 1-D intensity profile

$$P(t) = I(x_i - tN_x, y_i - tN_y) \quad (20)$$

at an offset of $t = \delta$ pixels, where $\delta > 0$.

In the case of extrapolation, we would like to evaluate $P(t)$ at $t = 0$ (i.e., at a boundary point), but the intensity at the boundary is unreliable due to the occlusion. This value can, however, be extrapolated from $P(t)$ with values $t > 0$. We assume that the intensity profile can be modeled with an exponential:

$$P(t) = \alpha t^\beta. \quad (21)$$

The model parameters, α and β , are determined using least-squares estimation² on $\log(P(t))$. In our results, we consider $P(t)$ for $t = 1, \dots, 15$, for this estimation. The intensity at the boundary, $P(0)$, is then simply determined by evaluating Equation (21) at $t = 0$. This entire process is repeated for each point along the occluding boundary. For objects of constant reflectance across the entire object, the extrapolation method is desirable, as it yields more accurate intensity estimates.

²The model parameters in Equation (21) are determined using least-squares estimation on $\log(P(t))$. Specifically, in the log domain we have $\log(P(t)) = \log(\alpha) + \beta \log(t)$. A quadratic error function in the model unknowns then takes the form $E(\vec{v}) = \left\| M\vec{v} - \vec{b} \right\|^2$, where each row of the matrix M is $[1 \ \log(t_i)]$, each corresponding entry of the vector \vec{b} is $[\log(P(t_i))]$, and $\vec{v} = [\log(\alpha) \ \beta]^T$. The least-squares estimate of this error function is $\vec{v} = (M^T M)^{-1} M^T \vec{b}$.

In the results of Section 3.1 (infinite light source), the method of simply measuring the intensity near the boundary was employed – measurements were made 1 pixel from the boundary. In the results of Section 3.2 and 3.3 (local and multiple light source), the extrapolation technique was employed. The reason for this difference is that the objects in our local and multiple light source experiments consisted of spheres of constant reflectance, which lend themselves to the extrapolation method. On the other hand, the objects in our infinite light source experiments did not have constant reflectance across the entire object, making it unlikely that the extrapolation method would yield accurate results.

In all cases, we converted the original color (RGB) image to grayscale ($\text{gray} = 0.299R + 0.587G + 0.114B$) from which the intensity measurements were made.

Finally, the values of λ in the error functions of Equation (10) and (18) were empirically determined to be 10 (infinite light source), and 1 (local light source). These values were held fixed for all examples given in the next three sections.

3.1 Infinite Light Source

Shown in Figures 4 and 5 are eight images of objects illuminated by the sun on a clear day. In order to determine the accuracy of our approach, a calibration target, consisting of a flat surface with a rod extending from the center, was placed in each scene. The target was approximately parallel to the image plane, so that the shadow cast by the rod indicated the direction of the sun. Errors in our estimated light source direction are given relative to this orientation.

The average estimation error is 4.8° with a minimum and maximum error of 0.6° and 10.9° . The image returning the largest error, the parking meters, is shown in the bottom panel of Figure 4. There are probably at least three reasons for this error, and for errors in general. The first is that the metallic surface violates the Lambertian assumption. The second is that the paint on the meter is worn in several spots causing the reflectance to vary, at times, significantly from patch to patch. And the third is that we did not calibrate the camera so as to remove luminance non-linearities (e.g., gamma correction) in the image.

Shown in the top panel of Figure 3 is an image of John Kerry and Jane Fonda sharing a stage at an antiwar rally. This image was circulated in February of 2004 in an attempt to discredit John Kerry during his campaign for the U.S. presidency. Shortly after its release, however, this image was determined to be a fake, having been created by digitally compositing two separate images. Although we do not know the true illuminant direction, we found an inconsistency in the estimated light direction: 123° for Kerry and 86° for Fonda. Note that these estimates were made on a very low quality image, giving hope that our algorithm will be fairly robust to, among others, JPEG compression. Shown in the bottom panel of Figure 3 is an authentic, although perhaps unlikely, image of Richard Nixon and Elvis Presley. The estimated light directions for each person are consistent, with Nixon at 98° and Presley at 93° .

3.2 Local Light Source

Shown in Figure 6 is an example of our experimental setup for testing the local light source estimation. In the top three panels are images of a pair of spheres imaged with the light source at 3 of 34 different positions. The light consisted of



Figure 3: Shown above is a known forgery of John Kerry and Jane Fonda sharing a stage at an antiwar rally. The estimated light direction for Kerry is 123° , while the direction for Fonda is 86° . Shown below is an authentic image of Richard Nixon and Elvis Presley. The estimated directions for Nixon and Presley are 98° and 93° .

a lamp with an exposed bulb, and the room was otherwise devoid of any light. With the pair of spheres being placed on either side of the origin of a world coordinate system, the light was placed at 93 cm or 124 cm from the origin along the y -axis, and moved from -123 to $+123$ cm along the x -axis, in 17 steps. Shown in the lower panel of Figure 6 is a schematic of these light positions (squares), with the triangles representing the estimated light position. On average, the position of the light source is estimated within 11.2 cm, with a minimum and maximum error of 0.7 and 22.3 cm. These values correspond to an average, minimum, and maximum error of 9.0%, 0.4%, and 18% as a percentage of the distance from the light source to the origin. With respect to estimating just the orientation of the light source, the average error is 0.4° with a minimum and maximum error of 0.04° and 1.1° .

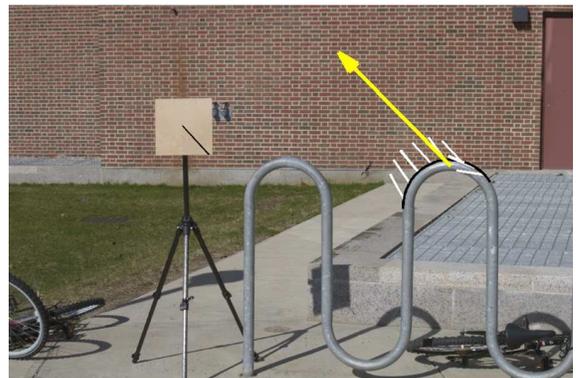


Figure 4: Shown are four images with the extracted occluding boundaries (black), individual light source estimates (white), and the final average light source direction (yellow arrow). In each image, the cast shadow on the calibration target indicates the direction to the illuminating sun, and has been darkened to enhance visibility.

Figure 5: Shown are four images with the extracted occluding boundaries (black), individual light source estimates (white), and the final average light source direction (yellow arrow). In each image, the cast shadow on the calibration target indicates the direction to the illuminating sun, and has been darkened to enhance visibility.

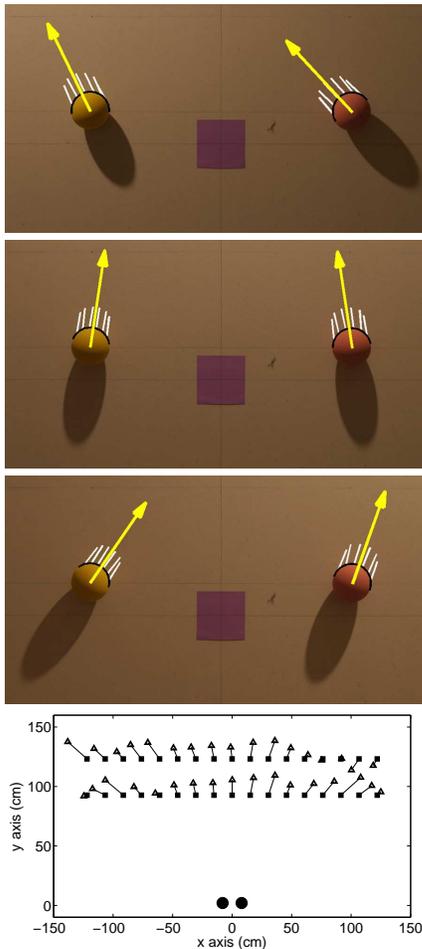


Figure 6: Shown in the top three panels are images of two spheres photographed with a local light source, with the extracted boundaries (black), individual light source estimates (white), and the final average light source direction (yellow arrow) for each sphere. Shown below is a schematic of our complete experimental setup along with the actual (squares) and estimated (triangles) light source positions.

3.3 Multiple Light Sources

Shown in Figure 7 are three synthetically generated images. In each case, the 3-D scene consisted of a single sphere illuminated with one or two infinite point light sources. In the two left-most panels the sphere was illuminated with a single light source positioned at -20° and $+20^\circ$ from vertical (90°). In the right-most panel, the sphere was illuminated with two lights positioned at $\pm 20^\circ$. Shown in each panel is the final estimated light source direction (yellow arrow). The actual light sources positions of the individual light sources are 70° and 110° yielding a virtual light source at 90° for the scene illuminated by both of these lights. The estimated directions are 69° , 108° , and 88° , yielding an average error of 1.7° . Note that in the case of the sphere illuminated with two light sources, the estimated direction is, as expected, the vector sum of the individual light sources.

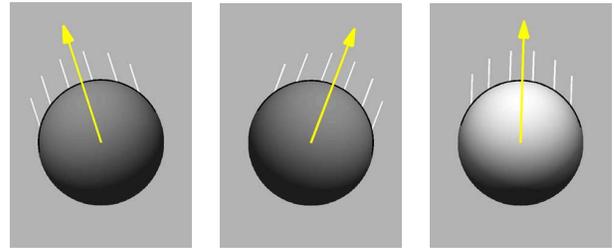


Figure 7: Shown are, from left to right, a sphere illuminated with a single light source positioned at -20° and $+20^\circ$ from vertical, and with two light sources positioned at $\pm 20^\circ$. Note that in the latter case, the estimate of the light source direction (yellow arrow) corresponds to the vector sum of the individual light source directions.

4. DISCUSSION

The creation of a digital forgery often involves combining objects/people from separate images. In so doing, it is difficult to exactly match the lighting effects due to directional lighting (e.g., the sun on a clear day). At least one reason for this is that such a manipulation may require the creation or removal of shadows and lighting gradients. And while large inconsistencies in lighting direction may be fairly obvious, there is evidence from the human psychophysics literature that human subjects are surprisingly insensitive to differences in lighting across an image [13, 7]. To the extent that the direction of the light source can be estimated for different objects/people in an image, inconsistencies in lighting can be used as evidence of digital tampering.

We have described a technique for estimating the direction (within one degree of freedom) of an illuminating light source. This technique builds upon work first described in [6]. We extended this basic formulation by relaxing some of the simplifying assumptions that were necessary to make the problem tractable, and by generalizing the approach to work under a local light source (e.g., floor lamp). We showed the efficacy of this technique on synthetically generated and real images.

We are currently investigating how surfaces of known geometry in the image (plane, sphere, cylinder, etc.) can be used to estimate the third component of the light source direction, N_z . If successful, this approach will remove the current ambiguity in the light source estimation. We are also investigating a technique to automatically determine which model, infinite or local, best describes the underlying image content so that a forensic analyst does not have to decide which model to use, as is currently the case.

We expect that the technique described here, in conjunction with other forensic tools (e.g., [4, 5, 10, 11]), will make it increasingly harder (but never impossible) to create convincing digital forgeries.

5. ACKNOWLEDGMENTS

This work was supported by an Alfred P. Sloan Fellowship, an NSF CAREER Award (IIS99-83806), a gift from Adobe Systems, Inc. and under a grant (2000-DT-CX-K001) from the U.S. Department of Homeland Security, Science and Technology Directorate (points of view in this document are those of the authors and do not necessarily represent the official position of the U.S. Department of Homeland Security or the Science and Technology Directorate).

Appendix A

This section derives the regularization term, Equation (17), for the local light source model of Section 2.3.

Consider the local light source estimates from a pair of objects (or a pair of boundaries from a single object or a pair of local patches from a single boundary) estimated by minimizing the quadratic error function of Equation (7). Denote these estimates as \vec{L}_1 and \vec{L}_2 , and denote \vec{c}_1 and \vec{c}_2 as the center pixel along each boundary. Assuming that these estimates are not parallel, the intersection of the individual light sources is determined by solving

$$\vec{c}_1 + \alpha_1 \vec{L}_1 = \vec{c}_2 + \alpha_2 \vec{L}_2 \quad (22)$$

for the scalar values α_1 and α_2 , using standard least-squares estimation. This intersection yields an estimate of the position of the local light source, $\vec{L} = \vec{c}_1 + \alpha_1 \vec{L}_1$.

Consider now the collection of individual estimates along each patch of an occluding boundary, $\vec{L}^1, \dots, \vec{L}^n$. Under the model of a single local light source, each of these estimates should be in the direction $\vec{L} - \vec{c}_i$, where \vec{c}_i is the center pixel of patch i . The regularization term, therefore, penalizes each estimate \vec{L}^i proportional to its deviation from this direction. Specifically, the penalty is proportional to the difference between the initial estimate \vec{L}^i and the projection of the estimate onto $\vec{L} - \vec{c}_i$:

$$\begin{aligned} \vec{R}_i &= \vec{L}^i - \vec{\Delta}_i \left(\vec{\Delta}_i^T \vec{L}^i \right) \\ &= \left(I - \vec{\Delta}_i \vec{\Delta}_i^T \right) \vec{L}^i \\ &= C_i \vec{L}^i, \end{aligned} \quad (23)$$

where I is the identity matrix and where,

$$\vec{\Delta}_i = \frac{\vec{L} - \vec{c}_i}{\left\| \vec{L} - \vec{c}_i \right\|}. \quad (24)$$

The penalty for \vec{L}^i is then simply the magnitude of \vec{R}_i , Equation (17). Note that \vec{L} , and hence the matrix C_i , is re-estimated on each iteration of the conjugate gradient minimization.

Appendix B

This section describes the minimization of the error function in Equation (18) (local light source model) using the conjugate gradient method.

Conjugate gradient minimization is an iterative technique for finding the minimum of a continuous function. Shown in Figure 8 is pseudocode for the version of this minimization technique that was employed in this paper [12].

The minimization of a continuous function $E(\vec{v})$ begins at a point \vec{v}_0 , and searches along a direction $\vec{\Delta}$ for a point \vec{v}_1

CONJUGATE-GRADIENT(E, \vec{v})

```

1   $i \leftarrow 0$ 
2   $k \leftarrow 0$ 
3   $\vec{r} \leftarrow -E'(\vec{v})$ 
4   $\vec{\Delta} \leftarrow \vec{r}$ 
5   $\delta_{new} \leftarrow \vec{r}^T \vec{r}$ 
6   $\delta_0 \leftarrow \delta_{new}$ 
7  while  $i < i_{max}$  and  $\delta_{new} > \epsilon^2 \delta_0$ 
8      do  $\delta_u \leftarrow \vec{\Delta}^T \vec{\Delta}$ 
9           $\alpha \leftarrow -\frac{E'(\vec{v})^T \vec{\Delta}}{\vec{\Delta}^T E''(\vec{v}) \vec{\Delta}}$ 
10          $\vec{v} \leftarrow \vec{v} + \alpha \vec{\Delta}$ 
11          $j \leftarrow 0$ 
12         while  $j < j_{max}$  and  $\alpha^2 \delta_u > \epsilon^2$ 
13             do  $\alpha \leftarrow -\frac{E'(\vec{v})^T \vec{\Delta}}{\vec{\Delta}^T E''(\vec{v}) \vec{\Delta}}$ 
14                  $\vec{v} \leftarrow \vec{v} + \alpha \vec{\Delta}$ 
15                  $j \leftarrow j + 1$ 
16          $\vec{r} \leftarrow -E'(\vec{v})$ 
17          $\delta_{old} \leftarrow \delta_{new}$ 
18          $\delta_{new} \leftarrow \vec{r}^T \vec{r}$ 
19          $\beta \leftarrow \delta_{new} / \delta_{old}$ 
20          $\vec{\Delta} \leftarrow \vec{r} + \beta \vec{\Delta}$ 
21          $k \leftarrow k + 1$ 
22         if  $k = m$  or  $\vec{r}^T \vec{\Delta} \leq 0 \quad \triangleright m = \dim(\vec{v})$ 
23             do  $\vec{\Delta} \leftarrow \vec{r}$ 
24                  $k \leftarrow 0$ 
25          $i \leftarrow i + 1$ 
26 return  $\vec{v}$ 

```

Figure 8: Pseudocode for conjugate gradient minimization – see Appendix B.

such that $E(\vec{v}_1) < E(\vec{v}_0)$. This search direction is opposite the direction of the gradient of $E(\vec{v})$ at \vec{v}_0 . At each iteration, the process is repeated with the search proceeding from the previous stopping point. The process terminates when a maximum number of iterations, i_{max} , has been reached, or if on the i^{th} iteration the gradient is below a tolerance ϵ . The initial point, \vec{v}_0 , is determined from the least-squares solution of $E_1(\cdot)$, Equation (7). Described next is the computation of the required gradient, $E'(\cdot)$, and Hessian, $E''(\cdot)$.

The error function of Equation (18) is composed of two terms:

$$E_1(\vec{v}) = \left\| M\vec{v} - \vec{b} \right\|^2, \quad (25)$$

and the regularization term:

$$E_2(\vec{L}^1, \dots, \vec{L}^n) = \sum_{i=1}^n \left\| C_i \vec{L}^i \right\|^2, \quad (26)$$

where the matrix M is given by Equation (8), the vector \vec{v} contains the individual light estimates \vec{L}^i and the ambient term A given in Equation (7), the vector \vec{b} is given in Equation (7), and the matrix C_i is given in Equation (23). The error function $E_2(\cdot)$ may be written in a more compact and convenient form as:

$$E_2(\vec{v}) = \|C\vec{v}\|^2, \quad (27)$$

where the block-diagonal matrix C is:

$$C = \begin{pmatrix} C_1 & & & 0 \\ & C_2 & & 0 \\ & & \ddots & \vdots \\ & & & C_n & 0 \end{pmatrix}, \quad (28)$$

and the form of each matrix C_i is described in Appendix A. The error function of Equation (18) then takes the form:

$$E(\vec{v}) = \|M\vec{v} - \vec{b}\|^2 + \lambda\|C\vec{v}\|^2, \quad (29)$$

and is differentiated to yield the gradient:

$$E'(\vec{v}) = 2M^T M\vec{v} - 2M^T \vec{b} + 2\lambda C^T C\vec{v}, \quad (30)$$

and twice differentiated to yield the Hessian:

$$E''(\vec{v}) = 2M^T M + 2\lambda C^T C. \quad (31)$$

Note that the matrix C is recomputed at every iteration of the minimization (i.e., it depends on the estimate of each \vec{L}^i at each iteration).

6. REFERENCES

- [1] R. Brunelli. Estimation of pose and illuminant direction for face processing. *Image and Vision Computing*, 15(10):741–748, 1997.
- [2] W. Chojnacki and M. J. Brooks. Revisiting Pentland’s estimator of light source direction. *Journal of the Optical Society of America*, 11(1):118–124, 1994.
- [3] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley Publishing Company, Inc., 2nd edition, 1993.
- [4] J. Fridrich, D. Soukal, and J. Lukáš. Detection of copy-move forgery in digital images. In *Proceedings of DFRWS*, 2003.
- [5] J. Lukáš, J. Fridrich, and M. Goljan. Determining digital image origin using sensor imperfections. In *Proceedings of the SPIE*, volume 5685, pages 249–260, 2005.
- [6] P. Nillius and J.-O. Eklundh. Automatic estimation of the projected light source direction. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2001.
- [7] Y. Ostrovsky, P. Cavanagh, and P. Sinha. Perceiving illumination inconsistencies in scenes. Technical Report AI Memo 2001-029, Massachusetts Institute of Technology, 2001.
- [8] A. Pentland. Finding the illuminant direction. *Journal of the Optical Society of America*, 72(4):448–455, 1982.
- [9] J.-M. Pinel, H. Nicolas, and C. L. Bris. Estimation of 2D illuminant direction and shadow segmentation in natural video sequences. In *Proceedings of VLBV*, pages 197–202, 2001.
- [10] A. C. Popescu and H. Farid. Exposing digital forgeries by detecting traces of resampling. *IEEE Transactions on Signal Processing*, 53(2):758–767, 2005.
- [11] A. C. Popescu and H. Farid. Exposing digital forgeries in color filter array interpolated images. *IEEE Transactions on Signal Processing*, (in press), 2005.
- [12] J. R. Shewchuck. An introduction to the conjugate gradient method without the agonizing pain. Technical Report CMU-CS-94-125, Carnegie Mellon University, 1994.
- [13] P. Sinha. Perceiving illumination inconsistencies. In *Investigative Ophthalmology and Visual Science*, volume 41/4:1192, 2000.