UC Berkeley
UC Berkeley Previously Published Works

Title

Exposing photo manipulation with inconsistent shadows

Permalink

https://escholarship.org/uc/item/7p63c7nk

Journal

ACM Transactions on Graphics, 32(3)

ISSN

0730-0301

Authors

Kee, Eric
O'Brien, James F
Farid, Hany

Publication Date

2013-06-01

DOI

10.1145/2487228.2487236

Supplemental Material

https://escholarship.org/uc/item/7p63c7nk#supplemental

Peer reviewed

Exposing Photo Manipulation with Inconsistent Shadows

Eric Kee, Dartmouth College

James O'Brien, University of California, Berkeley

and

Hany Farid, Dartmouth College

We describe a geometric technique to detect physically inconsistent arrangements of shadows in an image. This technique combines multiple constraints from cast and attached shadows to constrain the projected location of a point light source. The consistency of the shadows is posed as a linear programming problem. A feasible solution indicates that the collection of shadows is physically plausible, while a failure to find a solution provides evidence of photo tampering.

Categories and Subject Descriptors: I.2.10 [**Artificial Intelligence**]: Vision and Scene Understanding—*Scene Analysis*; I.3.6 [**Computer Graphics**]: Methodology and Techniques—*Image Forensics*; K.4.m [**Computing Milieux**]: Computers and Society—*Miscellaneous*

Additional Key Words and Phrases: Shadows, image forensics, photo manipulation, image manipulation, forgery detection

ACM Reference Format:

1. INTRODUCTION

Recent advances in computational photography, computer vision, and computer graphics allow for the creation of visually compelling photographic fakes. The resulting undermining of trust in photographs impacts law enforcement, national security, the media, advertising, e-commerce, and more. The nascent field of photo forensics has emerged to help restore some trust in digital photographs [Farid 2009].

In the absence of any digital watermarks or signatures, techniques in image forensics operate on the assumption that most forms of tampering will disturb some statistical or geometric property of an image. In a well-executed forgery these disturbances will either be perceptibly insignificant or they may be noticeable
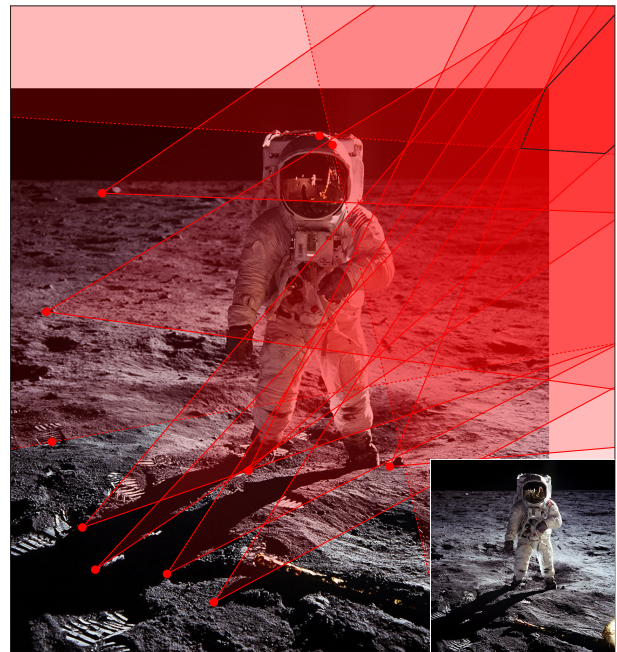
Original image copyright 1969, NASA

Fig. 1. Our algorithm finds that the shadows in this 1969 moon landing photo are physically consistent with a single light source. The solid lines correspond to constraints from cast shadows and dashed lines correspond to constraints from attached shadows. The region outlined in black, which extends beyond the figure boundary, contains the projected light locations that satisfy all of these constraints.

but subjectively plausible. Methods for forensic analysis provide a means to detect and quantify specific types of tampering. To the extent that these perturbations can be quantified and detected, they can be used to objectively invalidate a photo.

We describe a technique for determining if cast and attached shadows in a photo are consistent with the model of a single distant or local point light source. Forensic techniques based on analyzing lighting and shadows are attractive because 3-D lighting effects can be difficult to modify using commercial photo editing software, and low quality images can be analyzed since lighting effects and shadows survive common operations such as image compression and down-sizing.

There is some evidence that the visual system is capable of detecting small changes in lighting direction in simple controlled settings [Koenderink et al. 2004; Mamassian 2004; Khang et al. 2006; Pont and Koenderink 2007; Koenderink et al. 2007; O'Shea et al. 2010]. In more complex settings, however, the visual system is far less capable at detecting gross inconsistencies in lighting [Jacobson and Werner 2004; Ostrovsky et al. 2005; Farid and Bravo 2010]. In a forensic setting, a multitude and variety of cast and attached shad-

ows from complex shapes are cast onto equally complex and varied surfaces. Such limitations of the visual system imply that a forger may overlook inconsistencies in lighting and shadowing, and a visual inspection of shadows will, at best, be highly subjective.

The iconic photo of the 1969 moon landing, Fig. 1, provides an example of the variety and complexity of shadows that are common in photos. In fact, it has been argued by conspiracy theorists that the shadows in this photo are physically implausible and hence evidence of photo tampering and broader nefarious conspiracies. Beyond a subjective visual analysis, the physical consistency of shadows can be objectively determined by considering their basic geometry.

Consider a ray that connects a point in a shadowed region to its corresponding point on the shadow-casting object. In the 3-D scene, this ray intersects the light source. In a 2-D image of the scene created under linear perspective, the projection of this ray remains a straight line that must connect the images of the shadow point and object point, and intersect the projected image of the light source. These constraints hold regardless of the geometry of the object and the surface onto which the shadows are cast, and for either an infinitely distant or local light. Multiple constraints can therefore be used to determine the projected location of a light source in the image plane. Note that this projected location corresponds to an infinite number of 3-D light positions. We concern ourselves with the 2-D projection because a single image typically does not allow one to compute the 3-D location of the light in the scene.

If a scene purportedly contains a single light source but the shadows in the scene specify mutually inconsistent constraints that cannot be satisfied by any single light position, then this inconsistency evidences photo tampering. It can sometimes be difficult or impossible to precisely match a point on a shadow to its corresponding point on an object — particularly on attached shadows that form as a surface smoothly curves to face away from the light. We therefore consider a relaxed, conservative constraint in which the location of points on the object are restricted to a range of possible locations. These relaxed constraints specify either angular wedges or half-planes in the image that restrict the projected location of the light source. (See Fig. 2.) The consistency of multiple such constraints is framed as a linear programming problem. A viable solution is interpreted to mean that the shadows are physically plausible while a failure to find a solution is used as evidence of photo tampering.

1.1 Related Work

Previous lighting-based forensics methods estimate the 2-D lighting direction or lighting environment from the shading on an object's contour [Johnson and Farid 2005; 2007]. If the 3-D geometry of an object is known, then the 3-D lighting direction or lighting environment can be estimated [Kee and Farid 2010]. Related computer vision techniques that estimate lighting from a single image use object shading [Nillius and Eklundh 2001], or shadows cast onto planar surfaces [Sato et al. 2003; Okabe et al. 2004; Lalonde et al. 2011]. In [Karsch et al. 2011], manually approximated scene geometry is used to fit a local lighting model that is perceptually plausible, but insufficient for forensic application because the physical accuracy is heavily influenced by user input.

Photometric inconsistencies of a cast shadow's umbra were used to detect inconsistent shadows in [Liu et al. 2011]. Inconsistencies in the location of a cast shadow were used in [Zhang et al. 2009], but they placed several assumptions on the scene geometry: shadows were cast onto a planar ground plane and the objects casting shadows were vertical relative to the ground plane. In the most closely related work [Stork and Johnson 2006], the consis-

Original image copyright 2011, Geico Insurance

Fig. 2. Shown are: (top) a frame from the Geico commercial "Dunk – Easier Way to Save" depicting a somewhat incredible athletic performance; (middle) examples of cast (1-2) and attached (3) shadow constraints. The projected location of the light source must lie in the intersection of these constraints. Shown in the bottom is a multitude of constraints (solid lines correspond to cast shadows and dashed lines correspond to attached shadows). The shadows from the people and house are consistent with a light source located somewhere in the black-outlined region. The boy's shadow, however, is inconsistent with the rest of the scene.

tency of cast shadows in artworks was determined by identifying points on a cast shadow and their corresponding point on an object,
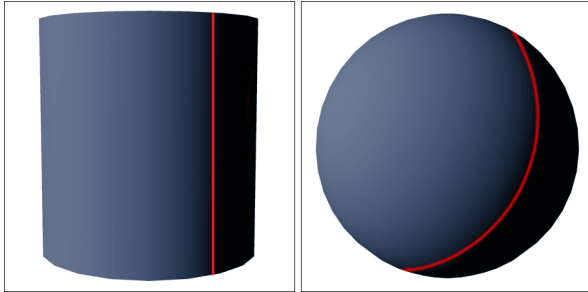
Fig. 3. A cylinder and sphere illuminated from the left with a distant point light. The red line and curve correspond to the terminator – the contour at which the surface normal is oriented 90 degrees from the direction to the light source. Points to the right of the terminator are in an attached shadow.

and then determining if these were consistent with a single light source. Our work expands this basic idea by relaxing the requirement that a strict shadow to object correspondence must be identified. We therefore allow a broader range of ambiguous cast shadows to be considered, including attached shadows. Unlike some previous shadow-based forensic techniques, we place no assumptions on the scene geometry.

2. METHODS

We first describe the geometry of cast and attached shadows and how these shadows restrict the projected location of a light source. Throughout, we assume a single distant or local light source and place no assumptions on the objects being illuminated or the surfaces onto which shadows are cast. We then frame the problem of determining if all shadows are consistent as a linear programming problem. In the case when the shadows are not consistent, we describe a simple randomized algorithm for finding an approximately minimal set of conflicting constraints that identifies the inconsistent shadows. We also describe a simple user interface that aids a forensic analyst in specifying constraints for cast and attached shadows.

2.1 Cast shadows

Shown in Fig. 2 is a frame from a Geico commercial with several shadows. Wedge-shaped constraints are used to describe the directions from a point in shadow to points on an object that may have cast the shadow. The wedge labeled 1 corresponds to a cast shadow on the roof. The wedge is fairly narrow because the tip of the cast shadow can reliably be determined to correspond to a point on the dormer. The wedge labeled 2 corresponds to a cast shadow on the garage roof. This second constraint is a wider wedge because the correspondence between the roof edge and its cast shadow is ambiguous.

The projected location of the light source should lie within the intersection of these wedges (which although illustrated as finite, are infinite in their extent beyond the figure boundary). Note that the wedges are oriented from the shadow towards the corresponding object. If, however, the light is behind the camera then these wedge constraints should be flipped 180 degrees about the selected shadow point. That is, due to perspective geometry there is a sign ambiguity as to the location of the projected light source (more on this issue below).
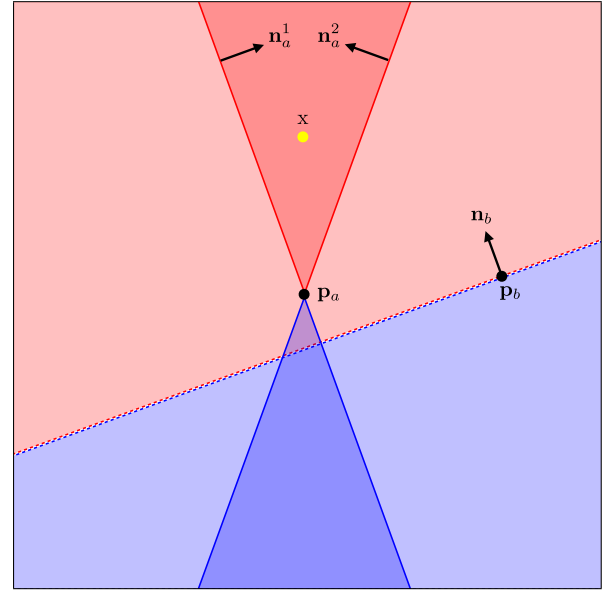

Fig. 4. A wedge constraint defined by two lines (solid red) and a half-plane constraint defined by one line (dashed red). The projected location of the light source (yellow dot, $\mathbf{x}$) lies within the region (outline in black) formed by the intersection of these constraints. A sign inversion of these constraints is depicted in blue.

2.2 Attached shadows

Attached shadows occur when objects occlude the light from themselves, leaving a portion of the object in shadow. For example, shown in Fig. 3 is a cylinder and sphere illuminated by a distant point light source positioned to the left. Points are in shadow if the surface normal makes an angle greater than 90 degrees with the direction toward the light. The terminator is defined to be the surface contour whose normals form a 90 degree angle with the direction toward the light source, as depicted by the red line and curve in Fig. 3.

Similar to a cast shadow, there is a correspondence between points in and out of shadow on either side of the terminator. This correspondence, however, can only be specified to within a half-plane because the light's elevation is ambiguous to within 180 degrees.

Although cylinders and spheres may not be particularly common in natural photographs, any locally convex surface can provide an attached shadow half-plane constraint. Shown in Fig. 2, for example, is an attached shadow constraint on a fold of the man's shirt (constraint 3 in the middle panel). Folds and other locally convex geometry are common and provide easily recognizable attached shadow constraints.

2.3 Forensics from shadows

For an authentic image there must be a location in the infinite plane ($\mathbb{R}^2$) that satisfies all cast and attached shadow constraints. That is, the intersection of all the constraints should define a non-empty region. In the bottom panel of Fig. 2, for example, there is a non-empty intersection for all of the shadow constraints on the parents and house — the region outlined in black. Note, however, that the shadow of the airborne boy in the background generates a wedge constraint that does not intersect the other constraints and is therefore physically inconsistent for a scene with one light source.

Cast and attached shadow constraints can be represented as linear inequalities in the plane. The satisfiability of these constraints can then be determined using standard linear programming. Shown in Fig. 4 are two lines defined implicitly by their normals $\mathbf{n}_a^1$ and $\mathbf{n}_a^2$, and the point $\mathbf{p}_a$. The direction of the normals specifies the region in the plane in which the solution $\mathbf{x}$ must lie. The intersection of these two regions is the upward facing wedge (red). Also shown in Fig. 4 is a single line defined by its normal $\mathbf{n}_b$ and point $\mathbf{p}_b$. This line specifies a half plane constraint in which the solution must lie. In each case, a shadow constraint is specified by either a pair of lines (wedge cast shadow constraint) or a single line (half-plane attached shadow constraint).

Formally, a half-plane constraint is specified with a single linear inequality in the unknown $\mathbf{x}$:

$$\mathbf{n}_i \cdot \mathbf{x} - \mathbf{n}_i \cdot \mathbf{p}_i \geq 0, \tag{1}$$

where $\mathbf{n}_i$ is normal to the line and $\mathbf{p}_i$ is a point on the line. A wedge-shaped constraint is specified with two linear constraints:

$$\mathbf{n}_i^1 \cdot \mathbf{x} - \mathbf{n}_i^1 \cdot \mathbf{p}_i \geq 0 \quad \text{and} \quad \mathbf{n}_i^2 \cdot \mathbf{x} - \mathbf{n}_i^2 \cdot \mathbf{p}_i \geq 0. \tag{2}$$

A collection of half-plane and wedge constrains can be combined into a single system of m inequalities:

$$\begin{pmatrix} \mathbf{n}_1 \\ \mathbf{n}_2 \\ \vdots \\ \mathbf{n}_m \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} - \begin{pmatrix} \mathbf{n}_1 \cdot \mathbf{p}_1 \\ \mathbf{n}_2 \cdot \mathbf{p}_2 \\ \vdots \\ \mathbf{n}_m \cdot \mathbf{p}_m \end{pmatrix} \geq \mathbf{0} \tag{3}$$

$$\mathbf{N}\mathbf{x} - \mathbf{P} \geq \mathbf{0}. \tag{4}$$

Given error-free constraints from a consistent scene, a solution to this system of inequalities should always exist. However we can account for errors or inconsistency by introducing a set of m slack variables s_i:

$$\mathbf{N}\mathbf{x} - \mathbf{P} \geq -\mathbf{s} \tag{5}$$
$$\mathbf{s} \geq \mathbf{0}, \tag{6}$$

where s_i is the i^{th} component of the m-vector $\mathbf{s}$. If the constraints are fully satisfiable then a solution will exist where all slack variables are zero. Solutions with a non-zero slack variable s_i mean that the solution is not consistent with constraint i.

The above inequalities can be combined into a single system as follows:

$$\begin{pmatrix} \mathbf{N} & \mathbf{I} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{s} \end{pmatrix} - \begin{pmatrix} \mathbf{P} \\ \mathbf{0} \end{pmatrix} \geq \mathbf{0}, \tag{7}$$

where $\mathbf{I}$ is an $m \times m$ identity matrix. We, of course, seek a solution that minimizes the amount of slack required to satisfy this system. As such, the linear program consists of minimizing the L_1 norm of the vector of slack variables,

$$\begin{pmatrix} \mathbf{0} & \mathbf{1} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{s} \end{pmatrix}, \tag{8}$$

subject to the constraints in Equation (7). That is, we seek a solution that minimizes the slack variables, while satisfying all of the cast and attached shadow constraints. If the slack variables for the optimal solution are all zero, then there exists a light position that satisfies all of the specified constraints. Otherwise, one or more of the shadows is inconsistent with the rest of the scene.

Recall that a shadow constraint is specified by connecting a point on a shadow with a range of possible corresponding points on an object. As described earlier there is an inherent sign ambiguity in specifying these constraints: if the light is behind the center of projection, then its projected location in the image plane is inverted and the constraint normals should all be negated. In this case, the constraints simply take the form:

$$-\mathbf{N}\mathbf{x} + \mathbf{P} \geq -s. \tag{9}$$

In Fig. 4, for example, the downwards facing region shaded blue corresponds to a sign inversion of the upwards facing region shaded red. In practice we solve both linear programs (with constraints Equation (5) and (9)) and select the solution with the minimal L_1 norm, Equation (8). If either the regular or inverted system has a solution with zero slack, then we conclude that the constraints are mutually consistent. Otherwise, there is no light position that is consistent with all of the constraints and we conclude that some of the constraints are being generated by parts of the image that have been manipulated.

When an image generates inconsistent constraints, we may wish to know which constraints are in conflict with others. These conflicting constraints provide the essential evidence that can be used to invalidate a forgery, and can be useful in determining what parts of an image may have been manipulated. We greedily find an approximately minimal set of inconsistent constraints.[1] To begin, two constraints are selected at random. If these constraints are not satisfiable, then they form a minimal set of inconsistent constraints. If they are satisfiable, then a randomly selected constraint is added to the set and the linear program is solved. Constraints are added in this way until the system is no longer satisfiable. This entire process is repeated with different random starting conditions. The smallest set of violating constraints, which may or may not be unique, provides a succinct summary of which parts of an image may have been altered.

2.4 User Interface

A forensic analyst's specification of cast shadow constraints is made easier and more efficient with a simple user interface. An analyst specifies cast shadow constraints by selecting a point in shadow and selecting a wedge region that comfortably encompasses the corresponding object casting the shadow, Fig. 5(a). After the first constraint is entered, a yellow wedge-shaped region is automatically rendered at the user's cursor, Fig. 5(b). This yellow region is the smallest wedge, emanating from the current point, that would completely cover both the forward and backward feasible regions of the linear programs associated with the previous constraint(s). The analyst's goal is to reduce the size of this wedge by identifying new constraints that more tightly encompass the shadow-casting object, Fig. 5(c). If the wedge that the analyst would need to specify in order to encompass the shadow-casting object fully encloses the yellow wedge, then the new constraint is redundant with previous ones and can simply be ignored, Fig. 5(d). If the wedge specified by the analyst does not overlap the yellow wedge at all, then the new constraint is inconsistent with previous ones, indicating a forgery.

For attached shadow constraints, specification is facilitated by automatically computing a shadow's orientation at a user specified point. Shown in Fig. 5(e), for example, is a point on an attached shadow (red dot) and the automatically rendered terminator orientation (pink line) and direction to the bright side of the shadow (pink arrow). The orientation and direction are determined by simply computing the local intensity gradient in a small spatial neighborhood. The analyst can either accept the specified shadow con-

---

[1]We use an approximation algorithm because finding the guaranteed minimal set of inconsistent constraints requires worst case exponential time.

Fig. 5. A user interface for specifying shadow constraints: (a) two cast shadows are manually selected on the roof; (b-c) an analyst specifies a point on a cast shadow (red dot) and automatically the range of plausible light directions is rendered (yellow wedge). Specification of a constraint (c) reduces the width of the wedge to conform to the object casting the shadow (red lines); (d) a constraint may be ignored if the yellow wedge is too narrow to reliably encompass the object which cast the shadow; (e-f) an analyst specifies a point on an attached shadow (red dot) and the orientation and direction of the shadow terminator is automatically suggested by computing the image gradient (pink dotted line and arrow) after which an analyst can accept or adjust the constraint.

straint (Fig. 5(f)) or rotate it slightly if necessary. Useful attached shadow constraints are found when the half-plane identified by the attached shadow excludes a portion of the yellow region (which does not occur in Fig. 5(e)).

This simple interface allows an analyst to efficiently identify what is typically a rich set of cast and attached shadow constraints. Shown in Fig. 6 is a computer generated scene which illustrates

the variety of shadows that an analyst might encounter. When cast shadows are well defined (e.g., the cones), the analyst will typically specify a single wedge, and when cast shadows are ambiguous (e.g., the spheres), the analyst will typically specify multiple wedges that combine to provide a tighter constraint than can be described by a single wedge. The analyst's selection of multiple wedges is guided by the yellow region that is visualized by the
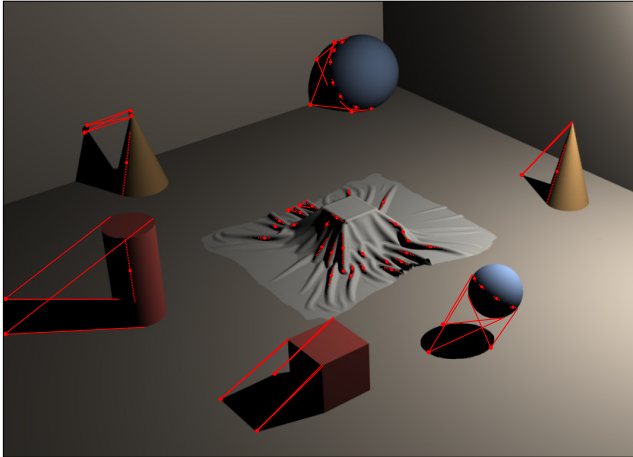
Fig. 6.   A computer generated scene with sixty cast and attached shadows. Wedge constraints from the cast shadows are shown with solid lines. Half-plane constraints from attached shadows are shown with dashed lines.
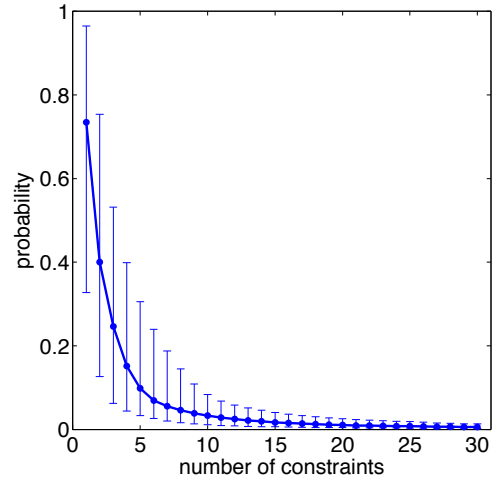

Fig. 7.   The median probability that a randomly selected light position will satisfy a variable number of cast and attached shadow constraints. The error bars correspond to the 25^{th} and 75^{th} quantile.
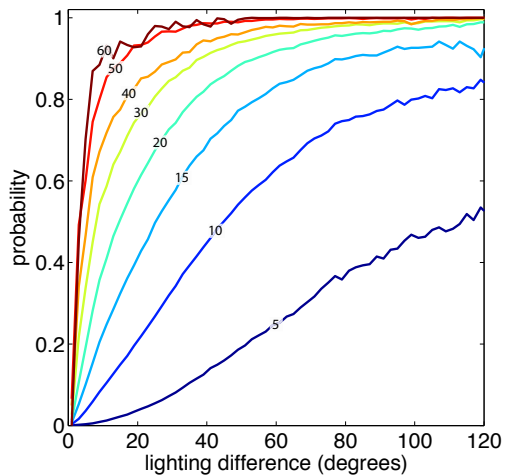

Fig. 8.   The probability of detecting a composite photo as a function of the angular deviation in the position of the light. Each curve corresponds to a different number of shadow constraints.

graphical interface, and wedges are typically located on opposing sides of an ambiguous cast shadow. When specifying attached shadow constraints, as on the cloth, a multitude of somewhat redundant shadows may be present, and an expeditious analyst will identify and omit these shadows by using the yellow region as a guide.

3.   RESULTS

We first validate our technique on a set of large-scale simulations and then provide results from several real-world examples of authentic and visually compelling forgeries.

3.1   Simulations

Shown in Fig. 6 is a computer generated scene with a variety of cast and attached shadows. This scene was rendered multiple times with either an infinitely distant or local point light in one of 49 or 98 locations, respectively. The distant lights were uniformly sampled over a hemisphere excluding an elevation less than 40 degrees. The local lights were uniformly sampled on the same hemisphere and at two different radial distances (R) from the ground plane ($R = 145$ and $R = 265$, where the height of the cylinder in the scene is 30).

We first assess how many wedge and half-plane shadow constraints are necessary to reasonably restrict the space of viable light positions. For each of the 147 computer generated scenes, exemplified by Fig. 6, a random subset of n cast and attached shadow constraints were first extracted[2]. Then, a random 3-D light position was drawn from a hemisphere, excluding an elevation less than 10 degrees, and projected into the image plane. Shown in Fig. 7 is the median probability that this light position was consistent with n constraints (horizontal axis) — the error bars correspond to the 25^{th} and 75^{th} quantile. The difference between the distant and local light are insignificant and so the results are combined.

With only one constraint, the median probability is 0.73 meaning that a single constraint from forged parts of the image is unlikely

to be effective in detecting photo composites. However, this probability decreases quickly as more constraints are added: with three constraints the probability is 0.25, with five constraints the probability is 0.10, with eight constraints the probability is 0.05, and with twenty constraints the probability is 0.01.

We next simulated the creation of a total of 1,008,714 composite photos with inconsistent lighting. Each composite consisted of the seven objects shown in Fig. 6. A varying sized subset of these objects (94 distinct subsets in total) were illuminated with the light source in one of 147 positions, and the remaining objects were illuminated with a different light source, for a total of 10,731 light pairings. Once composited, a maximum of 1,500 random sets of between 5 and 60 shadow constraints were extracted from each scene to yield approximately 150 million different systems of constraints. These constraints were selected so that the ratio between the wedge cast shadow and the half-plane attached shadow constraints was on average 1:2 (this is the same ratio of constraints found in the orig-

---

[2]A subset of n cast and attached shadow constraints were selected from a scene with N constraints as follows. If the total number of subsets of size n was less than 500, then all subsets were considered, otherwise, a random selection of 500 of all possible $N!/n!(N - n)!$ subsets were considered.
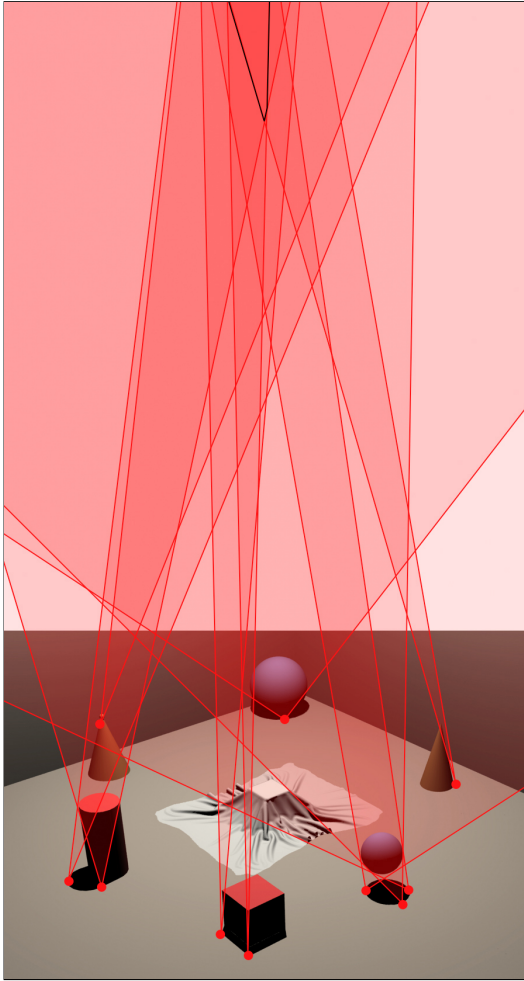
Fig. 9. Shown is an authentic photo and the successful result of a consistency check. The region of plausible locations for the light source is outlined in black.
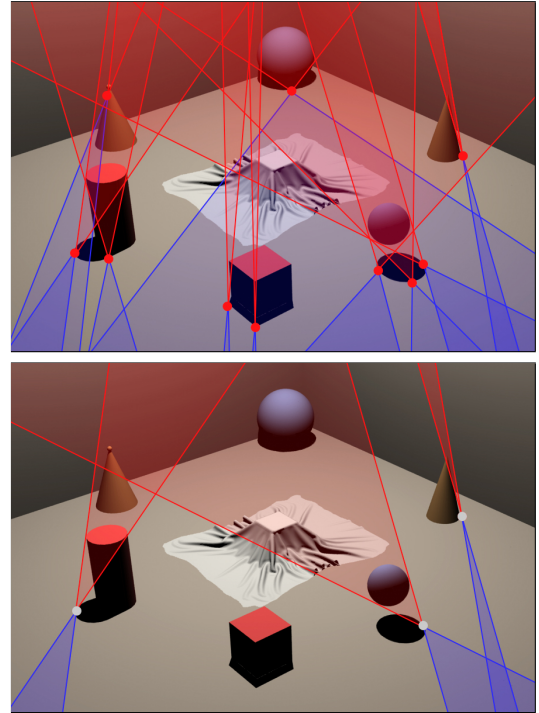

Fig. 10. Shown is a fake scene in which the floating sphere's shadow is inconsistent with the rest of the scene. Shown in the upper panel are the complete set of user specified constraints. Shown in the lower panel is a minimal set of violating constraints automatically determined from the above set of constraints. The red shaded regions corresponds to the positive constraints and the blue shaded regions correspond to the negative constraints.

inal scenes). If no solution is found that satisfies the selected constraints, then the photo is correctly classified as a fake.

Shown in Fig. 8 is the accuracy with which a forgery was detected in these scenes. The individual curves correspond to the total number of constraints. The horizontal axis corresponds to the angle between the two projected light directions in the composite photo, which is computed as the median angle between the projected directions of the two light sources at each constraint. The detection accuracy improves with the total number of constraints because more constraints carve out an increasingly smaller valid solution space.

With only 10 constraints, the difference in the lighting must reach $100°$ before the probability of detecting the fake reaches 80%. With 20 constraints, this same accuracy is achieved for a lighting difference of $35°$, and with 50 constraints even a small $10°$ discrepancy can be detected with a probability of 80%. At the same time, the accuracy of correctly classifying an authentic photo in these scenes is 100%. Because any portion of a cast or attached shadow can be used, it is fairly easy to find 50 or more constraints in a typical image.

Shown in panel of Fig. 9 is an authentic scene with eight constraints. The projected light position is in the intersection of the constraints which is outlined in black. Shown in Fig. 10 is a fake version of the scene shown in Fig. 9 in which the floating sphere from a differently lit scene was inserted. Shown in the upper panel are eight constraints selected in this scene and shown in the lower panel is the result from the automatic detection of a minimal set of unsatisfiable constraints. The red shading corresponds to the positive constraints and the blue shading corresponds to the negative constraints — neither are satisfiable. The median angular lighting difference at each constraint is $3.6°$ which in this case is enough of a difference to create an inconsistency in the shadows.

In the supplemental document that accompanies this paper we applied our method to several synthetic images from a previously published study [Farid and Bravo 2010]. In that study subjects were presented with fairly simple images where shadows were cast by two sets of objects onto two separate surfaces, and the subjects were asked to indicate if the shadows were mutually consistent (correct) or inconsistent. The results of that study showed that people performed poorly at this task even for these relatively simple scenes containing well-defined shapes with strong salient features. As shown in Sup. Fig. 1, our method works well on these images that confound human viewers.

3.2 Real World

Shown in Fig. 11 is an authentic photo illuminated with a distant light source (the sun). Also shown is the result of a successful consistency check for eleven cast shadows (dozens of other consistent
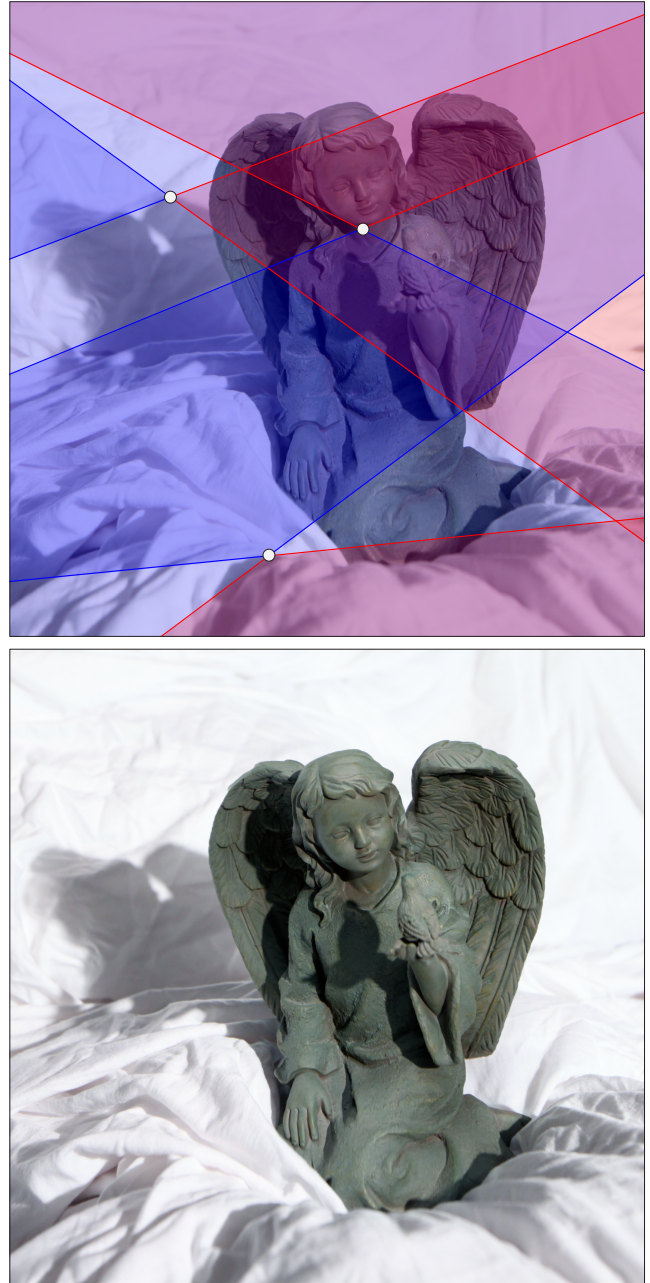
Fig. 11. An authentic photo and the result of a shadow consistency check. The region of plausible locations for the light source is outlined in black.

shadows are available in the scene but their constraints were not rendered to avoid clutter).

Fig. 12. A composite photo and a minimal set of violating constraints (these were automatically selected from a total of 31 user specified constraints). The shadows on the cloth cast by the angel's right wing and by another part of the cloth, and cast on the neck by her chin are mutually inconsistent. The red shaded regions corresponds to the positive constraints and the blue shaded regions correspond to the negative constraints.

Shown in Fig. 12 is a composite photo illuminated with a local light source where the shadows on the angel are inconsistent with the cloth and the angel's cast shadow. Also shown in this figure is a minimal set of violating constraints. A second composite along with a minimal set of violating constraints is shown in Fig. 15
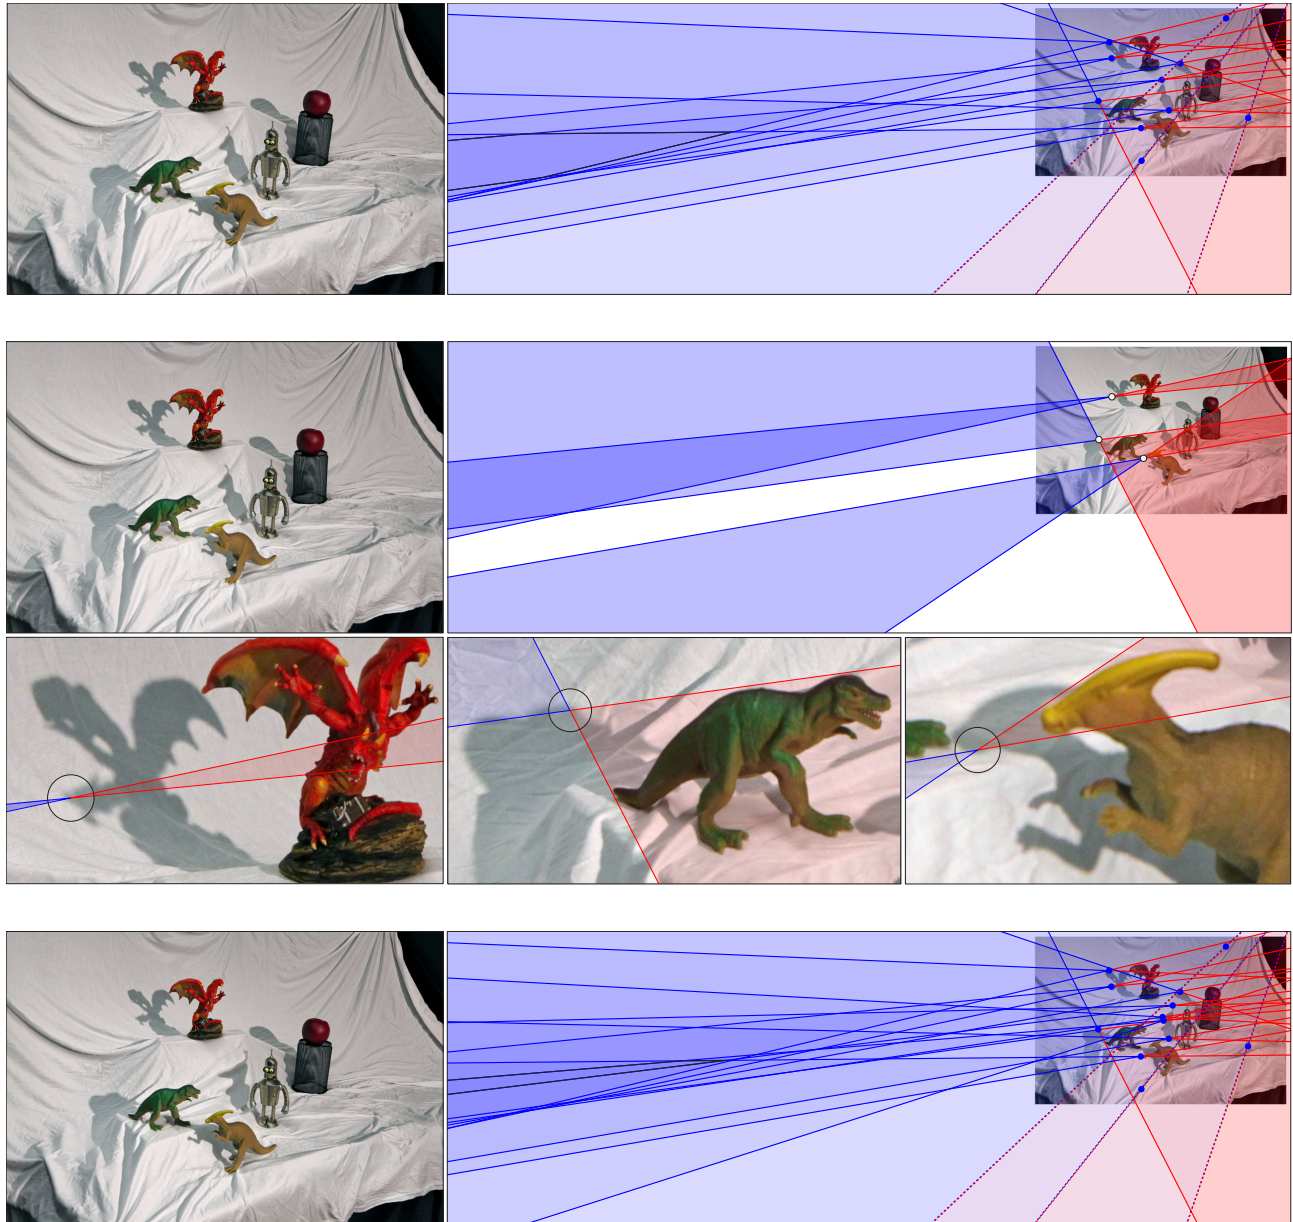
Fig. 13. An authentic photo (first row) and two fake photos (second and fourth row), and the results of a shadow consistency check. The solid lines correspond to cast shadows and dashed lines correspond to attached shadows. The red shaded regions correspond to the positive constraints and the blue shaded regions correspond to the negative constraints. Where a feasible solution is found, it is outlined in black. The panels in the third row show magnified views of the constraints from the image in the second row. For the authentic image in the top row a valid solution is found. For the fake image in the second row, the right panel shows a minimal set of constraints that are violated revealing that the duck-bill dinosaur's shadow is inconsistent with the rest of the scene (these were automatically determined from a total of 10 user specified constraints). The fake image in the bottom row illustrates a failure case for our method. The robot's shadow has been modified, but the modification still admits a valid solution.

Shown in the top row of Fig. 13 is a scene illuminated with a local light source. Also shown is the result of a successful consistency check for ten cast and attached shadows (many other consistent shadows are available in the scene but were not rendered to avoid clutter). Shown in the second row of Fig. 13 is an altered photo in which the duck-billed dinosaur and its shadow were inserted from a separate photo taken with the light in a different

position. The panel to the right shows a minimal set of mutually conflicting constraints revealing that the shadow of the duck-billed dinosaur is inconsistent with the rest of the scene. Although significantly different from the correct shadow, the shape and location of the dinosaur's shadow in the composite photo is not, at a casual inspection, obviously fake. Further, subjective arguments about the shape or direction of the shadow would be confounded by the irreg-

ular geometry upon which the shadow is cast. In addition to inserting the dinosaur, the image was also modified by rotating the apple and black container counter-clockwise while leaving their shadows unchanged. The apple's shadow, however, remains consistent by our test because the wedge specifying the cast shadow constraint is wide enough to encompass the intersection of the other constraints.

Shown in the bottom row of Fig. 13 is an altered photo that demonstrates a failure case for our method. The shadow of the robot was manipulated using Photoshop. However the manipulation purposefully shifted the shadow along rays from the light to the original shadow point. The result of this contrived manipulation is a distorted shadow that still generates constraints consistent with the true light location.

Shown in Fig. 1 is a shadow analysis of a photo from the Apollo 11 moon landing. The constraints correspond to four cast shadows from the astronaut's legs onto the moon's surface, three cast shadows from objects on the moon's surface, and three attached shadows on the astronaut's helmet, and footprint on the moon's surface. The region outlined in black identifies locations that satisfy all 11 of these constraints (dozens of other consistent shadows are available in the scene but were not rendered to avoid clutter). Thus we conclude that by our measure the lighting in this image is consistent.

3.3   Usability

We performed an informal survey of five users with some photo editing experience but no experience with our shadow analysis. After one hour of instruction and training, all five of the users successfully completed the analysis of the image shown in Fig. 13. The average amount of time taken was 30 minutes. Although the fifth user successfully completed the analysis, he employed one constraint that was invalid. In a post interview, he agreed that it was invalid, and was able to subsequently complete the analysis.

In certain cases, the specified shadow constraints will restrict the position of the projected light source to be in a closed polygonal region. Although not rendered as such, this is also the case in Figs. 1, 9, and 11. In these cases, it is more difficult for a forger to add a consistent shadow. In addition, this restriction to a finite area is a stronger constraint than previous work that identified inconsistencies in the projected 2-D light direction [Johnson and Farid 2005].

Lastly, we note that a local or infinite light source may sometimes project to a point at infinity in the image plane. A local light, for example, will project to infinity if the direction from the camera center to the light is parallel to the image plane. In this case, lines connecting points on a cast shadow to their corresponding points on an object will be parallel and hence intersect at infinity. Shown in Fig. 14 is such an example where the dashed lines connecting shadow and object on the cone and cube are parallel. An analyst limited by finite image resolution, however, cannot unambiguously identify exact points to specify these parallel lines and instead must specify a broader wedge shape (solid lines). These wedges diverge and therefore their intersection includes finite locations (black region) as well as the true light projection at infinity. Although lights projecting to infinity would create problems for algorithms attempting to find the light position by intersecting nearly parallel rays [Stork and Johnson 2006], this situation does not pose any particular difficulty in specifying a valid region using wedge and half-plane constraints.

4.   DISCUSSION

We have described a geometric technique for detecting photo manipulation based on inconsistent shadows. Previous lighting- and shadow-based forensic techniques exploited cast shadows for which a unique object-to-shadow pairing was available (e.g., the tip of a cone). In contrast, we exploit a broad range of ambiguous cast and attached shadows. Although each constraint is typically not highly informative, a combination of many such constraints can be highly effective in detecting inconsistent shadows that are not perceptually obvious. The subsequent determination of shadow consistency is framed as a standard linear programming problem affording a computationally efficient solution. A consistent solution to all specified shadow constraints is interpreted to mean that the shadows are physically plausible while a failure to find a solution is used as evidence of photo tampering. In the latter case, it can be difficult to visually identify which constraints are inconsistent. We, therefore, developed a method to identify the inconsistent constraints thus providing insight into which parts of an image were manipulated.

This method is intended for use where there is only a single dominant light source. While this limitation does preclude analysis of scenes lit by multiple point lights or diffuse area lighting, it includes the common situation of outdoor scenes lit by the sun or indoor scenes photographed with a flash. Scenes with multiple light sources or strong interreflections are usually evident by a corresponding multitude of shadows for a single object. It is likely that this basic approach will extend to area lights [Ramamoorthi and Hanrahan 2001] that yield reasonably well defined shadows.

Beyond an assumption of linear perspective projection and a single dominant point light source, no other assumptions about the scene geometry or photometry are required. Should lens distortion be an issue, standard techniques can be employed to estimate and remove lens distortion. Alternatively the constraint wedges could be expanded to accommodate the bounded movement of features in the image plane due to lens or other type of distortion.

A key step in applying our method is for the analyst to select a set of shadows from the image and to specify appropriate wedges or half-planes. A poor selection of constraints could, of course, lead to a failure in detecting a manipulated image. To minimize the likelihood of this, we described a simple user interface which makes the annotation of shadows relatively fast and easy to perform.

Even with this interface, our method relies on a user correctly selecting constraints. The strength of our approach is that these constraints can be objectively validated because the correspondence between objects and their shadows is generally clear. In cases where the relation is not clear, large encompassing wedges may be specified or other less ambiguous shadows may be used. Our method therefore shifts the dialogue from "does the lighting/shadow look correct?" (which is well known to be highly unreliable), to a discussion of whether an analyst has correctly selected the range of points on an object that correspond to a point on a shadow (a far more objective task). In this regard, our method lets humans do what computers are poor at — understanding scene content — and lets the computer do what humans are poor at — assessing the validity of geometric constraints.

A sufficiently informed forger (e.g., [Shesh et al. 2009]) could, of course, use this forensic technique to ensure that all shadows are consistent. One way to counter this is to combine this shadow analysis with other techniques for estimating lighting from a single image [Nillius and Eklundh 2001; Okabe et al. 2004; Lalonde et al. 2011]. This addition will make it more difficult, but never impossible, to create a consistent and visually compelling fake.
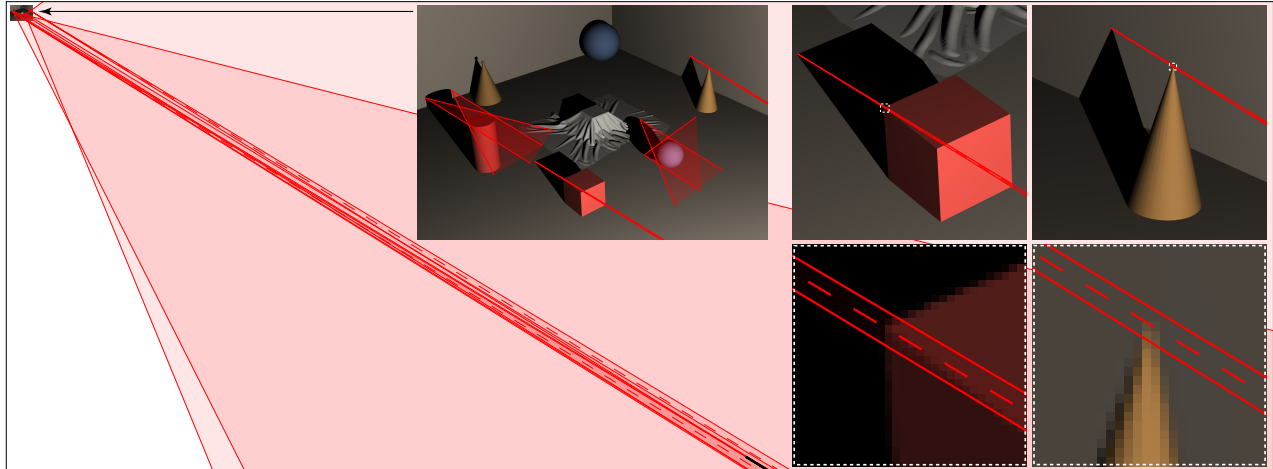
Fig. 14.   A scene is illuminated by a local light source that projects to a point at infinity in the image plane. In this special case, the light casts shadows that project to parallel lines in the image (dashed lines). A forensic analyst, limited by image resolution, must specify broader wedges which diverge (solid red lines), yielding an intersection that includes finite locations (black outline) as well as the true light projection at infinity.

ACKNOWLEDGMENTS

REFERENCES

FARID, H. 2009. A survey of image forgery detection. *IEEE Signal Processing Magazine 2,* 26, 16–25.

FARID, H. AND BRAVO, M. J. 2010. Image forensic analyses that elude the human visual system. In *SPIE Conference on Media Forensics and Security*.

JACOBSON, J. AND WERNER, S. 2004. Why cast shadows are expendable: Insensitivity of human observers and the inherent ambiguity of cast shadows in pictorial art. *Perception 33,* 11, 1369–1383.

JOHNSON, M. K. AND FARID, H. 2005. Exposing digital forgeries by detecting inconsistencies in lighting. In *ACM Multimedia and Security Workshop*. 1–10.

JOHNSON, M. K. AND FARID, H. 2007. Exposing digital forgeries in complex lighting environments. *IEEE Transactions on Information Forensics and Security 3,* 2, 450–461.

KARSCH, K., HEDAU, V., FORSYTH, D., AND HOIEM, D. 2011. Rendering synthetic objects into legacy photographs. In *SIGGRAPH Asia*. 157:1–157:12.

KEE, E. AND FARID, H. 2010. Exposing digital forgeries from 3-D lighting environments. In *Workshop on Information Forensics and Security*.

KHANG, B.-G., KOENDERINK, J. J., AND KAPPERS, A. 2006. Perception of illumination direction in images of 3-D convex objects: Influence of surface materials and light fields. *Perception 35,* 5, 625–645.

KOENDERINK, J. J., PONT, S., VAN DOORN, A., KAPPERS, A., AND TODD, J. 2007. The visual light field. *Perception 36,* 11, 1595–1610.

KOENDERINK, J. J., VAN DOORN, A., AND PONT, S. 2004. Light direction from shad(ow)ed random Gaussian surfaces. *Perception 33,* 12, 1405–1420.

LALONDE, J.-F., EFROS, A. A., AND NARASIMHAN, S. G. 2011. Estimating the natural illumination conditions from a single outdoor image. *International Journal of Computer Vision*.

LIU, Q., CAO, X., DENG, C., AND GUO, X. 2011. Identifying image composites through shadow matte consistency. *IEEE Transactions on Information Forensics and Security 6,* 3, 1111–1122.

MAMASSIAN, P. 2004. Impossible shadows and the shadow correspondence problem. *Perception 33,* 11, 1279–1290.

NILLIUS, P. AND EKLUNDH, J.-O. 2001. Automatic estimation of the projected light source direction. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. *Computer Vision and Pattern Recognition*.

OKABE, T., SATO, I., AND SATO, Y. 2004. Spherical harmonics vs. Haar wavelets: Basis for recovering illumination from cast shadows. *IEEE Conference on Computer Vision and Pattern Recognition 1*, 50–57.

O'SHEA, J. P., AGRAWALA, M., AND BANKS, M. S. 2010. The influence of shape cues on the perception of lighting directions. *Journal of Vision 12,* 10, 21:1–21.

OSTROVSKY, Y., CAVANAGH, P., AND SINHA, P. 2005. Perceiving illumination inconsistencies in scenes. *Perception 34*, 1301–1314.

PONT, S. AND KOENDERINK, J. 2007. Matching illumination of solid objects. *Attention, Perception, & Psychophysics 69*, 459–468.

RAMAMOORTHI, R. AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM Press, 497–500.

SATO, I., SATO, Y., AND IKEUCHI, K. 2003. Illumination from shadows. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 25,* 3 (march), 290 – 300.

SHESH, A., CRIMINISI, A., ROTHER, C., AND SMYTH, G. 2009. 3D-aware image editing for out of bounds photography. In *Proceedings of Graphics Interface*. 47–54.

STORK, D. G. AND JOHNSON, M. K. 2006. Estimating the location of illuminants in realist master paintings: Computer image analysis addresses a debate in art history of the Baroque. In *International Conference on Pattern Recognition*. 255–258.

ZHANG, W., CAO, X., ZHANG, J., ZHU, J., AND WANG, P. 2009. Detecting photographic composites using shadows. In *IEEE International Conference on Multimedia and Expo*. 1042–1045.

Fig. 15.   A composite photo and a minimal set of violating constraints (these were automatically determined from a total of 25 user specified constraints). The red shaded regions corresponds to the positive constraints and the blue shaded regions correspond to the negative constraints. The three middle panels are a magnified view of the selected shadows.