Exposing Photo Manipulation with Inconsistent Reflections

James F. O'Brien
University of California, Berkeley
and
Hany Farid
Dartmouth College

The advent of sophisticated photo editing software has made it increasingly easier to manipulate digital images. Often visual inspection cannot definitively distinguish the resulting forgeries from authentic photographs. In response, forensic techniques have emerged to detect geometric or statistical inconsistencies that result from specific forms of photo manipulation. In this paper we describe a new forensic technique that focuses on geometric inconsistencies that arise when fake reflections are inserted into a photograph or when a photograph containing reflections is manipulated. This analysis employs basic rules of reflective geometry and linear perspective projection, makes minimal assumptions about the scene geometry, and only requires the user to identify corresponding points on an object and its reflection. The analysis is also insensitive to common image editing operations such as resampling, color manipulations, and lossy compression. We demonstrate this technique with both visually plausible forgeries of our own creation and commercially produced forgeries.

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

Additional Key Words and Phrases: Reflections, image forensics, photo manipulation, mirrors, image manipulation, forgery detection, center of projection.

ACM Reference Format:

O'Brien, J. F., and Farid, H. 2012. Exposing Photo Manipulation with Inconsistent Reflections. ACM Trans. Graph. 31, 1, Article - (January 2012), 11 pages. DOI=TBD http://doi.acm.org/TBD

1. INTRODUCTION

Photographs can no longer be trusted. Forged images have appeared in tabloid magazines, main-stream media outlets, political attacks, scientific journals, and the hoaxes that land in our email in-boxes. These doctored photographs are appearing with growing frequency and sophistication, and even experts often cannot rely on visual inspection to distinguish authentic images from forgeries.

On the political front, for example, a photograph of Senator John Kerry and Jane Fonda sharing a stage at an anti-war rally emerged during the 2004 presidential primaries as Senator Kerry was campaigning for the Democratic nomination. The photograph, however, was a fake. The picture of Senator Kerry was from a speech in June of 1971, and the unrelated picture of Jane Fonda was from August 1972. The two photographs were composited together to give the impression that Senator Kerry shared the controversial anti-war views of activist Jane Fonda [Light 2004]. On the scientific front, in 2004 Professor Hwang Woo-Suk of Seoul National University and colleagues published what appeared to be ground-breaking advances in stem cell research. Evidence slowly emerged that these results were manipulated and/or fabricated. After months of controversy, Hwang retracted the paper and resigned his position at the University. An independent Korean panel investigating the accusations of fraud found that at least nine of the eleven customized stem cell colonies that Hwang had claimed to have made were fakes. Much of the evidence for those nine colonies, the panel said, involved doctored photographs of two other, authentic, colonies [Wade 2005]. Finally, in the news media, *The Economist* was criticized when it published, in June 2010, a cover photo showing a solitary President Obama on the Louisiana beach inspecting the BP oil spill. The photo was accompanied with the headline "The damage beyond the spill", eluding to potential political problems facing President Obama as a result of the oil spill. This photograph, however, had been altered to remove two other people standing alongside the President [Peters 2010].

In addition to the ethical, political, and legal implications raised by this lack of trust in photography, studies have shown that doctored photographs can alter our own memories of actual events [Wade et al. 2002; Garry and Wade 2005; Garry and Gerrie 2005; Sacchi et al. 2007]. In one such study participants were shown original and doctored photographs of memorable public events at which they were present (*e.g.*, the 2003 protests against

This work was supported by NSF grants IIS-0915462 and CNS-0708209, UC Lab Fees Research grant 09-LR-01-118889-OBRJ, Intel's Science and Technology Center for Visual Computing, and by gifts from Adobe Systems, Microsoft, NVIDIA, and Pixar.

Authors' email address: job@berkeley.edu, farid@cs.dartmouth.edu

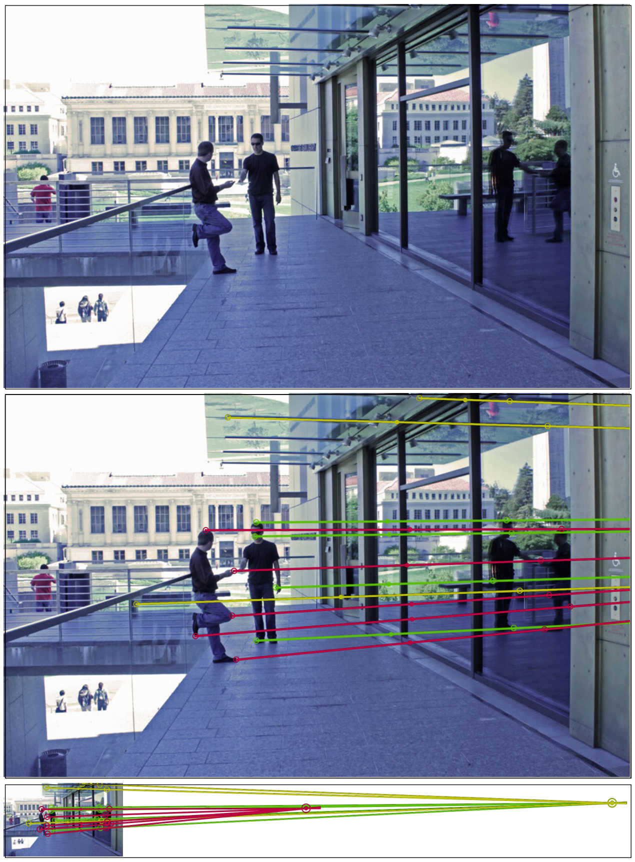

Fig. 1: *To most viewers the top image appears to be a genuine photograph. However, anaylsis of the relations between objects in the scene and their reflections (middle and bottom images) reveals inconsistencies and proves that the image has been manipulated. In this case, the figure leaning on the railing is inconsistent both with the standing figure and other scene elements. The green and yellow lines show that the standing figure, awning, and bench all share a single consistent reflection vanishing point (green open circle). The red lines show that the reflection vanishing point for the leaning figure (red open circle) is inconsistent with the rest of the scene.*

the Iraq war). The doctored photographs, showing either larger crowds or more violence, changed the way in which participants recalled the events [Sacchi et al. 2007]. This surprising finding is due to a number of factors including our natural trust in photographs and our general inability to easily detect doctored photographs.

High quality digital cameras and affordable photo-editing software such as Photoshop have made it easier for nearly anyone to create compelling photo fakes. In addition, recent advances in Computer Vision and Computer Graphics (*e.g.*, [Avidan and Shamir 2007; Hays and Efros 2007; Grabler et al. 2009; Ritschel et al. 2009; Sunkavalli et al. 2010]) point to a future of growing sophistication for photo manipulators. The need for forensic techniques for exposing photo fakery is, therefore, critical.

Digital watermarking has been proposed as a means for image authentication. The drawback of this approach is that a watermark must be inserted at the time of recording, which would limit this approach to specially equipped cameras. In contrast, passive forensic techniques operate in the absence of watermarks or signatures. These techniques work on the assumption that although photo manipulation may leave no obvious visual clues, it may alter some geometric or statistical property in the image. With this approach, a collection of complementary techniques are each designed to detect specific types of artifacts. The combination of many such techniques makes creating a fake photo more difficult (but not impossible) because the forger has to carefully consider a host of different possible artifacts that may not be visually obvious, and yet may be detectable by one of many forensic techniques. Even if a forger successfully fools most avenues of detection, a single failed test can conclusively and objectively discredit an image.

In this paper we describe one such forensic technique for detecting geometric inconsistencies in images containing reflections, such as the example shown in Figure 1. The underlying methodology is derived from basic rules of three-dimensional geometry, planar reflection, and linear perspective projection. This technique makes minimal assumptions about the scene geometry, assuming only that the reflecting surface is flat. The resulting forensic analysis for determining if a reflection is plausibly real or not relies on simple and intuitive computations, and a few minutes of user input. This forensic technique adds to a growing body of forensic analyses [Farid 2009] that, combined, promise to make it increasingly more difficult to create a compelling and undetectable forgery.

2. BACKGROUND

In many ways, the human visual system is remarkable, capable of hyperacuity [Westheimer and McKee 1975], rapid scene understanding [Potter 1976], and robust face recognition [Sinha et al. 2006]. In other ways, however, the visual system can be quite inept. For example, the visual system can be insensitive to inconsistencies in lighting and shadows [Ostrovsky et al. 2005], certain judgments of lightness and color [Adelson 2000], perspective distortions [Bravo and Farid 2001], viewing position [Vishwanath et al. 2005], shadows, and the geometry of reflections [Croucher et al. 2002; Bertamini and Parks 2005; Farid and Bravo 2010].

Photo manipulators often implicitly or explicitly take advantage of these insensitivities when altering photographs. As a result, it is relatively easy to create a convincing inauthentic photo with substantial physical and geometric inconsistencies [Kelby 2008]. For example, several techniques have been developed to create fake reflections that, although not geometrically correct, are visually plausible [Shelbourne 2007].

The past decade has seen the development of a number of forensic techniques for detecting a variety of different photo inconsistencies [Farid 2009]. These include techniques for detecting cloning [Fridrich et al. 2003; Popescu and Farid 2004; Pan and Lyu 2010], splicing [Ng and Chang 2004], resampling artifacts [Popescu and Farid 2005a; Kirchner and Gloe 2009], color filter array aberrations [Popescu and Farid 2005b; Kirchner 2010], disturbances of a camera's sensor noise pattern [Lukas et al. 2006; Fridrich 2009], chromatic aberrations [Johnson and Farid 2006; Gloe et al. 2010], and lighting inconsistencies [Johnson and Farid 2005; Johnson and Farid 2007b; Kee and Farid 2010; Riess and Angelopoulou 2010].

This paper focuses on geometric inconsistencies that arise when fake reflections are inserted into a photo or when a photo containing reflections is manipulated. Prior work addressed this issue but required an explicit estimate of the reflecting surface's three-dimensional normal [Farid and Bravo 2010]. By contrast, the imaged-based construction described here does not require a prior estimate of a surface normal and does not make any assumptions about the reflecting surface other than that it is flat. We also describe how noisy estimates of the image's center of projection can be used to identify faked images containing rectangular reflectors. An ad-

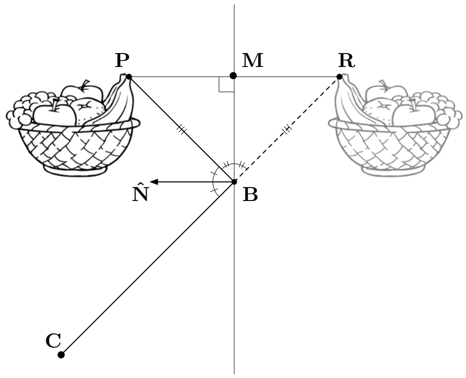

Fig. 2: *This figure illustrates the geometric relationships that support the existence of a reflection vanishing point for linear-perspective images of planar reflectors. In an image formed with center of projection* $\mathbf{C}$*, the line between a point,* $\mathbf{P}$*, on some object and the point,* $\mathbf{R}$*, on its reflection must be parallel to the surface normal,* $\hat{\mathbf{N}}$*, of the reflecting mirror. Because the points* $\mathbf{P}$*,* $\mathbf{B}$*,* $\mathbf{C}$*, and* $\mathbf{R}$ *must all be coplanar, this two-dimensional diagram illustrates the three-dimensional phenomena without loss of generality.*

ditional feature of our methods is that, unlike some other forensic techniques, they are insensitive to common operations such as resampling, color manipulations, and lossy compression.

Outside of the field of image forensics, others have also made use of the same geometric invariants used here. Artists, for example, use reflection vanishing points to draw perspective scenes containing mirrors [Montague 2010]. At the same time, artists, for the sake of composition, often deviate from the correct geometry of reflections [Bertamini et al. 2003; Cavanagh et al. 2008] and the methods described here may be of use to art historians wishing to analyze classical paintings.

Although planar reflecting surfaces do not appear with great frequency in natural environments (aside from lakes and other still bodies of water), many man-made environments feature mirrors, windows, computer screens, tabletops, and other flat specular surfaces that create reflections. Our method will only be applicable to images containing reflections, but this type of limitation is typical of photo forensic techniques that are designed to detect specific forms of manipulation. Combining a host of different and complementary forensic techniques provides a forensic analyst a powerful toolkit for disproving images.

We note that while a photo may be proven fake by detecting a statistical, lighting, or geometric inconsistency, the lack of any such inconsistency does not prove that a photo is authentic. Failure to detect an inconsistency only proves that it is either authentic or a very good fake. This one-sidedness is intrinsic to the nature of any type of image authentication scheme.

3. PLANAR REFLECTION GEOMETRY

Our objective is to develop tests that can be applied directly to an image in order to detect when the scene shown in the image has been manipulated in some fashion. To do so, we consider the underlying geometric relations that must occur in a photograph of a scene where one or more objects are visible both directly in the image and indirectly in a reflecting surface that appears in the image. By combining the three-dimensional geometric relations of planar reflection with the geometry of linear perspective projection we develop image-based tests that must invariably apply to legitimate photographs.

3.1 The Reflection Vanishing Point

Reflections are observed in a scene when light leaving some object bounces off of a specular surface and subsequently enters the aperture of the observing imaging device. For smooth planar reflectors, such as flat mirrors, window panes, or still water surfaces, the angle between the incoming direction and the surface normal is the same as that between the normal and the outgoing direction. Further, all three are coplanar. As shown in Figure 2, one can equivalently state by reciprocity that a ray $\mathbf{CR}$ from the observer, toward a point in the observed reflection will strike the reflecting surface at some point, and a ray $\mathbf{BP}$ from that point to the corresponding point on reflected object will be symmetric about the surface normal with the ray from the observer.

Although this basic description of reflection geometry is no doubt familiar to the reader, an often overlooked consequence of this geometry is that a ray between the reflected object and the apparent location of its reflection must appear to be perpendicular to the reflecting surface. To see why, consider that to the observer it appears as if the ray toward the reflecting surface continues through the surface and strikes a virtual object behind the mirror. The isosceles triangle formed by the points $\mathbf{R}$, $\mathbf{P}$, and $\mathbf{B}$ is bisected by the plane of the reflecting surface at $\mathbf{B}$ and therefore the line containing $\mathbf{R}$ and $\mathbf{P}$ must be perpendicular to that plane.

The above explanation places the virtual object formed by the reflection at a location such that $\mathbf{P}$ and $\mathbf{R}$ are equidistant from $\mathbf{B}$. For a single view, $\mathbf{R}$ could equally well be located anywhere on the ray from the observer through $\mathbf{B}$, and this distance uncertainty corresponds to an unknown scale factor for the virtual object. If multiple viewpoints were considered then the scale ambiguity would vanish and the only consistent location for all views would be the equidistant one. However, if only one image from a single viewpoint is being considered then the image of the line between $\mathbf{P}$ and the equidistant $\mathbf{R}$ is identical to the image of the line between $\mathbf{P}$ and any other possible location for $\mathbf{R}$ along the line containing $\mathbf{C}$ and $\mathbf{B}$. As a result it is still true that a ray between the reflected object and the apparent location of its reflection must appear in the image as if it were perpendicular to the reflecting surface in the three-dimensional scene.

The above reasoning applies for all points on reflected objects and it implies that for a given planar reflecting surface, the set of lines connecting points on objects to the corresponding points on their reflections will all appear as if they were rays in the three-dimensional scene that are perpendicular to the surface and therefore mutually parallel. Assuming that the image of the scene was created by linear perspective projection, the image of these parallel lines will form a bundle that converges to a common vanishing point, $\mathbf{v}$, as shown in Figure 3. In the special case where the reflecting surface is parallel to the view-plane normal then the lines will remain parallel in the image and the vanishing point for these parallel lines will be at infinity in the direction defined by the lines.

We refer to this vanishing point as the *reflection vanishing point* for a given reflecting surface. Lines in the image between a point on an object and its reflection in the surface must all converge to the same point. Further, this vanishing point is the vanishing point for lines perpendicular to the reflecting surface, and any scene elements that are known to be perpendicular to the reflecting surface must also be consistent with this vanishing point.

3.2 Reflection Line Midpoints

Another geometric implication of the isosceles triangle formed by the points $\mathbf{R}$, $\mathbf{P}$, and $\mathbf{B}$ is that the midpoint, $\mathbf{M}$, of the line segment $\mathbf{RP}$ must lie in the plane of the reflecting surface. (See Figure 2.)

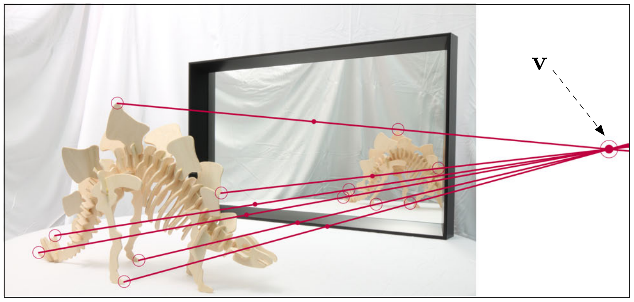

Fig. 3: *This photograph shows a wooden figure that can be viewed both directly and indirectly by its reflection in the planar mirror. The red lines connect features on the figure with corresponding features in its reflection. Because the image shows an unmanipulated reflection, the lines converge to a common vanishing point,* $\mathbf{v}$*, where the lines intersect. The points marked on each line between the features it connects show where the lines intersect the plane of the mirror.*

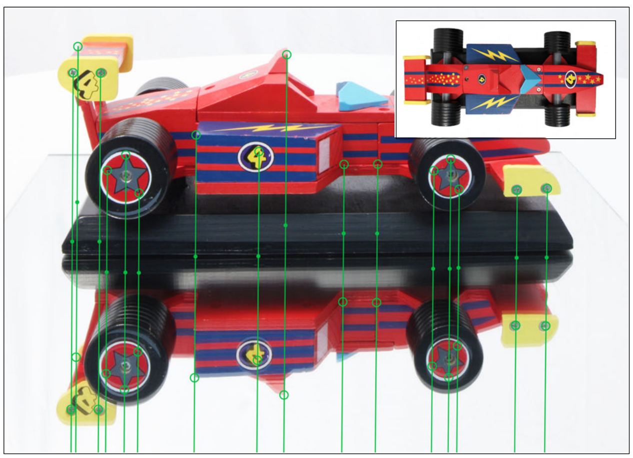

Fig. 4: *This image shows a toy car sitting on a reflective surface. The green lines connect features with their corresponding reflections, both of which are marked by hollow circles. The filled circles on each line mark the reflection line midpoints where the lines intersect the plane of the mirror.*

Computing these midpoints in the three-dimensional scene space is trivial, just take the average of $\mathbf{R}$ and $\mathbf{P}$. However, due to foreshortening the image of $\mathbf{M}$ will not be the midpoint between the images of $\mathbf{P}$ and $\mathbf{R}$. In other words, if $\mathbf{m}$, $\mathbf{p}$, and $\mathbf{r}$, are respectively the projections of $\mathbf{M}$, $\mathbf{P}$, and $\mathbf{R}$ into the image plane then in general $\mathbf{m} \neq (\mathbf{p}+\mathbf{r})/2$.

The image of the line between $\mathbf{P}$ and $\mathbf{R}$ is a degree-one rational polynomial that can be written as

$$\mathbf{l}(t) = \frac{\mathbf{x}_0 + \mathbf{x}_1 t}{h_0 + h_1 t} \tag{1}$$

where $\mathbf{x}_0$, $\mathbf{x}_1$, h_0, and h_1 are unknown. The numerator and denominator of Eq. (1) can be scaled arbitrarily without changing the result, so we remove this redundancy by scaling so that $h_1 = 1$. (In case where $h_1 = 0$ the line is parallel to the image plane, foreshortening is not an issue, and $\mathbf{m}$ is just the average of $\mathbf{p}$ and $\mathbf{r}$.)

Fixing the parametrization such that $t = 0$ and $t = 1$ respectively correspond to $\mathbf{p}$ and $\mathbf{r}$ produces

$$\mathbf{p} = \mathbf{l}(0) = \mathbf{x}_0/h_0 \quad \text{and} \tag{2}$$

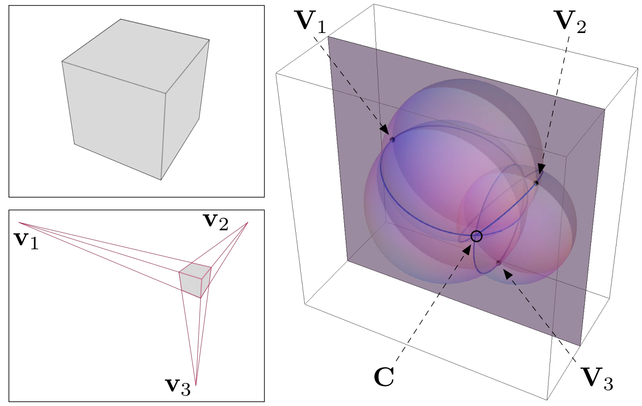

Fig. 5: *The image in the upper-left shows a perspective rendering of a cube where three vanishing points can be readily determined. The lower-left image has been expanded to show the three vanishing points in the two-dimensional image. The three-dimensional construction shown on the right illustrates how each pair of vanishing points defines a sphere and the intersection of these spheres determines the center of projection, with an ambiguity due to symmetry about the image plane.*

$$\mathbf{r} = \mathbf{l}(1) = (\mathbf{x}_0 + \mathbf{x}_1)/(h_0 + 1) \ . \tag{3}$$

The known location of the vanishing point implies that

$$\mathbf{v} = \lim_{t\to\infty} \mathbf{l}(t) = \mathbf{x}_1/h_1 = \mathbf{x}_1 \ . \tag{4}$$

Combining Eqs. (2)-(4) and solving for h_0 yields

$$h_0 = \frac{\mathbf{v}-\mathbf{r}}{\mathbf{r}-\mathbf{p}} \tag{5}$$

where either the x or y components of the vectors may be used to compute h_0 and will produce the same result because $\mathbf{v}$, $\mathbf{r}$, and $\mathbf{p}$ are collinear. (For numerical stability the component producing the largest divisor should be used.) The above equations can now be combined to compute the image midpoint as

$$\mathbf{m} = \mathbf{l}(1/2) = \frac{\mathbf{p}h_0 + \mathbf{v}/2}{h_0 + 1/2} \ . \tag{6}$$

Examples of these midpoints are shown in Figures 3 and 4. The points indicate the location of the intersection between the line connecting object points to their reflection and the plane containing the reflector. They can be compared to other scene elements for consistency. For example, in Figure 4 features that are farther in depth have reflection line midpoints higher up in the image, which is consistent with the reflecting plane that is pictured.

3.3 The Center of Projection

Although right angles occur rarely in natural scenes, like planar reflectors, they do appear quite commonly in man-made environments. Furthermore, one place where they regularly occur is in the frames around mirrors, window panes, and other flat specular surfaces. This situation is noteworthy because the edges of a rectangular frame provide two orthogonal directions for which vanishing points can be computed, and these directions are both orthogonal to the reflector's surface normal. When vanishing points for three mutually orthogonal directions are known, they can be used to compute the center of projection for an image [Kubovy 1986; Caprile and Torre 1990].

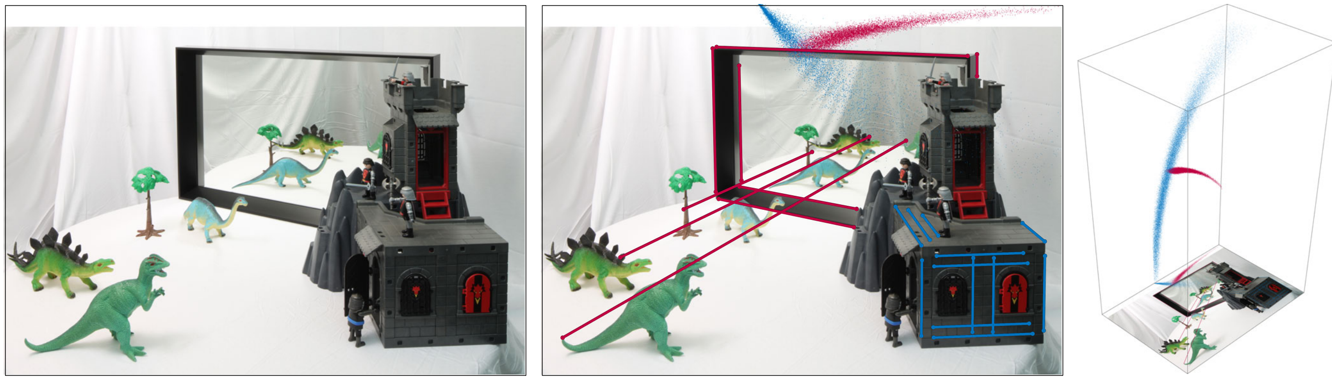

Fig. 6: *The image shown on the left contains both a number of dinosaurs whose reflections can be seen in a rectangular mirror, and also a toy castle with rectilinear sides. The edges of the mirror provide two orthogonal vanishing points that are orthogonal to the mirror's surface normal, and together they allow the center of projection to be computed. Similarly, the castle has three orthogonal sides which also allow the center of projection to be computed. Lines in the center image show features that were used for computing the vanishing points superimposed on the original image. Because the computation is unstable, random perturbation of feature points in a three-pixel diameter produces the clouds of possible centers shown in the middle and right images. The overlap between the two clouds shows that a plausible single center of projection exists that is consistent with both the dinosaur/mirror and castle.*

Let $\mathbf{v}_1$, $\mathbf{v}_2$, and $\mathbf{v}_3$ be vanishing points in an image for three mutually orthogonal directions, and let $\mathbf{V}_1$, $\mathbf{V}_2$, and $\mathbf{V}_3$ be the corresponding locations in a three-dimensional space where we have placed the image plane at $Z = 0$, as shown in Figure 5. If $\mathbf{C}$ is the center of projection then lines from $\mathbf{C}$ through each of the $\mathbf{V}_i$ will appear degenerate in the image, and so each of these lines must be parallel to all other lines that vanish at the respective $\mathbf{v}_i$. Because these three sets of lines are mutually orthogonal, a valid location for $\mathbf{C}$ must satisfy the following system of quadratic equations:

$$(\mathbf{C} - \mathbf{V}_1) \cdot (\mathbf{C} - \mathbf{V}_2) = 0 \quad (7)$$
$$(\mathbf{C} - \mathbf{V}_2) \cdot (\mathbf{C} - \mathbf{V}_3) = 0 \quad (8)$$
$$(\mathbf{C} - \mathbf{V}_3) \cdot (\mathbf{C} - \mathbf{V}_1) = 0 \quad . \quad (9)$$

Each of these equations is equivalent to requiring that $\mathbf{C}$ be on the surface of a sphere defined with two of the $\mathbf{V}_i$ as its poles. As shown in Figure 5 the solution, if one exists, is the pair of points where all three spheres intersect and this pair of points is symmetric about the image plane. The principal point, $\mathbf{c}$, is the projection of $\mathbf{C}$ into the image plane and it is well known to be the orthocenter of the $\mathbf{v}_i$ [Hartley and Zisserman 2004]. Appendix A discusses how the above system of equations can be used to solve for the three-dimensional center of projection which supplements the principal point with the distance to the image plane.

Unfortunately, the computation of either $\mathbf{C}$ or $\mathbf{c}$ is extremely unstable. The nature of this instability is readily apparent if one considers the two-dimensional analog: a pair of highly overlapping circles in the plane, where even small perturbations of the circles can substantially move the locations of the intersections. The problem is further exacerbated because the vanishing points themselves are computed from the intersection of lines in the image. These lines are parallel in the scene and their projections into the image are often nearly parallel so that computing their intersection is notoriously unstable as well. This situation means that computing $\mathbf{C}$ or $\mathbf{c}$ requires taking the result of one unstable computation and feeding it into another. Any approach that fails to consider this instability would be inherently flawed and unreliable.

We deal with this instability by attaching a measure of uncertainty to feature points that will be used to compute vanishing points and subsequently the center of projection. For example, a clearly imaged feature could have its position known within a two pixel wide region, but a blurry or partially obscured feature might be associated with a much larger region in the image. We then compute the centers of projection multiple times with random perturbations of the feature points over their region of uncertainty. As shown in Figure 6, the result is that rather than a single location for the center of projection, we instead have a cloud of plausible locations. Two sets of image features are consistent with each other if the resulting clouds overlap.

In determining overlap, it is important to note that the volume of the overlap does *not* correspond to an estimate of the likelihood of consistency. Even a small overlap volume implies that a single center of projection exists that could plausibly explain both sets of features. The fact that many other inconsistent locations exist is immaterial. To reinforce this fact we point out that in Figure 6 the overlap between the two clouds is quite small, but it nevertheless corresponds to the correct (and in this case known) location of the center of projection for this real, unmanipulated photograph.[1] We also note that image editing operations, such as cropping and uniform scaling, that move the center of projection relative to the image center do not disturb our test because we are comparing locations relative to the image features and not relative to the arbitrary coordinates of the image frame.

Figure 6 illustrates why the three-dimensional center of projection is more useful for comparison than the two-dimensional principal point. Even small amounts of uncertainty tend to smear the clouds out into streaks. In two dimensions these streaks are likely to overlap by coincidence. In three dimensions, the likelihood of coincidental overlap is greatly decreased and the test therefore is more discriminatory.

In addition to the technique used here for computing the center of projection from the vanishing points of three orthogonal directions, other methods exist for computing it based on other types of known geometry. For example, the center of projection can be com-

[1]The photograph was taken using a Canon TS-E 24mm f/3.5L II tilt-shift lens that allows the center of projection to be moved away from the center of the image.

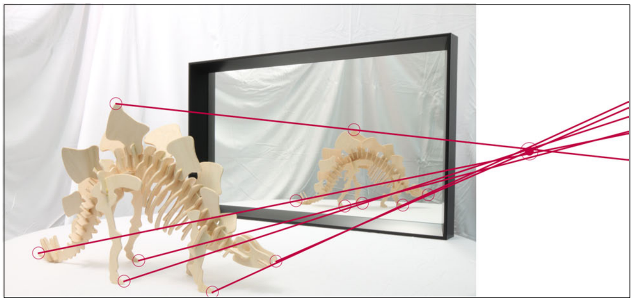

Image copyright James O'Brien and Hany Farid.

Fig. 7: *In this figure the image has been manipulated by composting the reflection from one photograph into another. The resulting image looks plausible, but the lack of a well-defined reflection vanishing point reveals it as a forgery.*

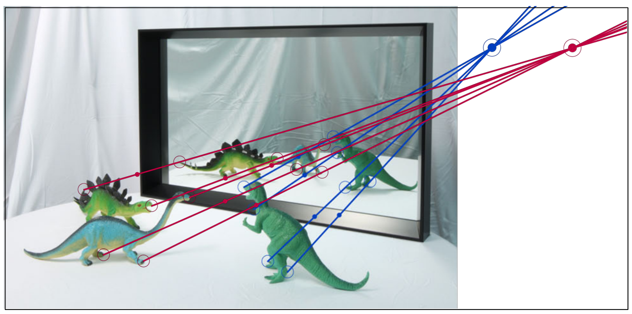

Image copyright James O'Brien and Hany Farid.

Fig. 8: *This composite image was formed by combining individual photographs where the mirror's position was varied. As a result, each object has a well-defined reflection vanishing point, but they are mutually inconsistent. The reflections of the herbivorous dinosaurs correspond to the pictured mirror location. The carnivorous dinosaur's reflection corresponds to a mirror location that is more rotated toward the viewer.*

puted from ellipses that are the images of spheres under perspective projection [Hartley and Zisserman 2004] and from images of human eyes [Johnson and Farid 2007a]. Regardless of the method, the center of projection should be consistent across all components of an image and the cloud-based approach we describe could be applied with any method for computing the center of projection.

4. DETECTING INCONSISTENCIES

The invariants described above provide a variety of tests that can potentially expose inconsistencies in an image. Naturally, not all of these tests are applicable to any given image and deciding which invariants apply and how they can be tested requires some human understanding of the image contents. As a result we do not have an automated algorithm that can be applied blindly to an image. Instead we have tests that require the insight of a human operator, yet are nevertheless objective and can conclusively detect some types of manipulation.

4.1 Ill-Defined Refection Vanishing Points

If a scene contains multiple features that appear in an image both directly and indirectly through a reflection, then those features should define a common reflection vanishing point. This construction is demonstrated in Figure 3 where all the lines between features and their reflections pass through a single well-defined point. In contrast, Figure 7 shows a composite image where the reflection has been copied from another photograph that had the wooden figure positioned differently.

Most viewers looking at the composite image have difficulty determining that it has been manipulated. Certainly a casual viewer would not notice the inconsistent reflection. Even when a viewer has decided that the reflection is inconsistent it remains difficult to construct an objective argument that could be used to convince another person beyond a reasonable doubt that the image has been manipulated. However, the failure of the image to produce a clean vanishing point unequivocally exposes the image as fake.

In a similar fashion, the image in Figure 8 appears realistic but is exposed as false by inconsistencies in the reflection vanishing points. In this case multiple objects can be viewed in a single reflector. While each object individually has a well-defined reflection vanishing point, different objects disagree on the location of this vanishing point. This disagreement violates the assumption that all three objects are being reflected by the same flat mirror, and it implies some form of manipulation.

4.2 Implausible Midpoints

The reflection line midpoints correspond to the projection of feature points into the plane of the mirror. In general it can be difficult to determine if a computed midpoint's location agrees with other scene elements. For example, in Figure 3 the midpoint for the topmost feature is placed in the center of the mirror's surface with no other features around it. It would be difficult to conclusively argue that some other point on that line could not be a valid midpoint without relying on some additional information such as the height of the wooden figure relative to the mirror.

However, in some instances it may be possible to reason about the scene geometry to arrive at assertions that can be tested. Again in Figure 3, the feet of the wooden figure are resting on the tabletop which appears perpendicular to the upright mirror. As a result, the lines between the feet and their reflection should be parallel to the tabletop so that the midpoints are approximately where the mirror plane and table top intersect. If the midpoints appeared elsewhere then we could conclude that the image must either be fake or that the mirror was not actually perpendicular to the table top.

While it may not always be possible to test the absolute position of a reflection line midpoint, the relative positions can also be useful. In both Figures 4 and 9 the nearly parallel lines to the reflection vanishing point indicate that the reflecting plane is nearly parallel to the view-plane normal. This configuration causes the midpoints to be sorted along the direction of the lines according to their distance from the image plane. In the unaltered image shown in Figure 4 the sorting correctly matches the depth of points on the car. However the building in Figure 9 has been inserted into the photograph and its midpoints appear out of order relative to other scene elements.

We also observe that if all midpoints for an object appear in a straight line then it implies both that the object must be a planar surface and that the surface is perpendicular to the mirror.[2] This situation would be rather uncommon for real scenes, but the commonly used photo-manipulation technique of pasting a reversed copy of an object back into a photograph produces a set of collinear midpoints and thus can be easily detected.

4.3 Inconsistent Centers of Projection

The method we have described for computing the center of projection using either a rectangular mirror or rectilinear structure pro-

[2]The other possibility is that the projection of the reflecting plane into the image is a line, but in that case the reflection would not be visible.

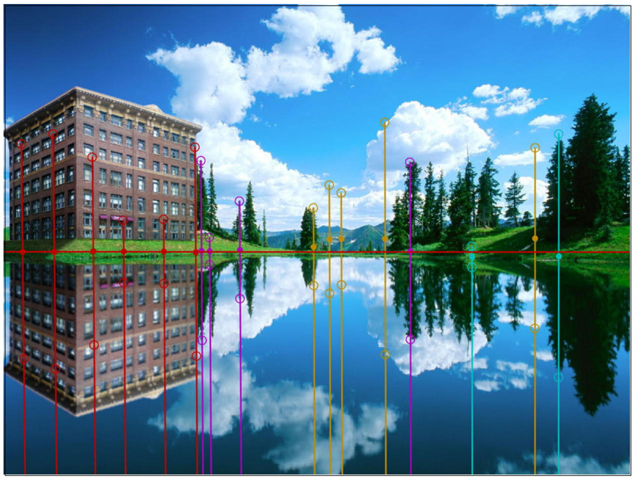

Original photo "Mirror Lake, Gunnison National Forest, Colorado," copyright 2008.

Fig. 9: *In this image a building has been composited into a natural scene. Because the plane of the lake recedes vertically from the viewer, the reflection line midpoints for more distant objects should appear higher in the image. The midpoints for features on the building, shown in red, violate this ordering and are inconsistent with other parts of the scene. In particular, the trees and rock immediately to the right of the building have midpoints, shown in purple, that indicate approximately the same depth as the building despite appearing behind it. It is also improbable that the two trees on the far right are closer than the building, yet their midpoints, shown in cyan, appear lower. Finally, midpoints for the building are all collinear which dubiously implies the same depth for all the building's features.*

duces clouds of possible locations. If two independent sets of image features produce non-overlapping clouds then this implies that no single center of projection could possibly have created the image. This situation is demonstrated in Figure 10 where the perspective has been shifted and the image of the castle has been scaled and repositioned. Because the castle's orientation and location in the scene are largely unchanged, a forensic analysis of lighting or shading would be unlikely to detect this manipulation. However, comparison of the center of projection clouds generated by the mirror and the castle shows that they do not overlap and reveals the presence of manipulation.

A similar example appears in Figure 11 where two different mirrors can each be used independently to compute the center of projection for the image. In the top row, a legitimate image produces overlapping clouds, but the image in the second row was built from two separate photographs. The first photograph was taken using an 18 mm focal length with the camera roughly two feet away from the subjects. The second photograph was taken with a 46 mm focal length but the camera was translated approximately six feet back along the view direction to compensate so that the arrangement of the subjects was almost identical. Even though the green dinosaur's situation in the image may appear unchanged to casual inspection, the longer focal length substantially alters the perspective convergence of parallel lines and produces a very different cloud of possible locations for the center of projection. As shown in the lower-right of Figure 11, the displacement of the blue cloud farther away from the image plane makes it immediately obvious that images with different focal lengths have been combined and the ratio of distances even provides information about the relative focal lengths. This example also demonstrates that the three-dimensional clouds are more discriminatory because although the three-dimensional clouds are clearly distinct, projection of the clouds into the image plane (*i.e.* the cloud of principle points) results in overlap.

The longer focal length also produces a cloud with larger variation. This behavior is expected because the smaller range of relative distances to the camera reduces the effects of perspective foreshortening, pushing the vanishing points out toward infinity, and decreasing the stability of the computation. It may be tempting to take the dramatically different variance of the two clouds as an indication of inconsistency, but doing so would be ill-advised because other factors, such as poorly separated feature points or surfaces nearly parallel to the image plane, can also produce changes in cloud variance.

4.4 Other Examples

The image in Figure 1 appears to depict a handoff between two people. Both men appear to be looking at the other and the more distant man appears to be handing an envelope to the closer one. However, this image was built from a composition of several photographs. In one photograph, the first man was holding a smart phone while waiting outside the building. After he left, another picture was taken with the second man pretending to hand an envelope to where the first person had been standing. Additional images were used to populate the background and complete the image. This image could plausibly be used as evidence that the first man had engaged in a transaction with the second, unsavory character. Conclusively convincing someone that the image is fictitious based on subjective visual inspection would be uncertain, but analysis of the reflection vanishing points shows that the reflection for the first man does not match the reflections of the building and second man. The double-paned windows create some uncertainty in the locations of reflection features, but not enough to accommodate a single reflection vanishing point. We also observe that additional reflections in the glass railing could be used to further discredit the image.

Figure 12 shows an unlikely situation where a giant cat is attacking a New York City street. While few people would credit this image as real, the image itself looks plausible and it is only the nonsensical content that makes it immediately suspect. We can nonetheless use this image to demonstrate our image analysis methods because the glass building-front contains prominent reflections. Several of the cars in the street present features that can be located in the reflection and these correspondences are connected with blue lines. In addition to these reflected features, several structures in the scene contain lines perpendicular to the reflecting surface. These structures include the building's awning, edges of the adjacent building, the crosswalk, and lines in the concrete sidewalk. In general these structures might not be perpendicular to the building, but in this instance, as shown by the green lines, they do agree with the reflection vanishing point for the cars.

The cat's reflection appears distorted with slight ripples and blur. This distortion combined with the rounded, fuzzy nature of the cat makes it difficult to definitively locate corresponding points on both the cat and its reflection. Regardless, we can still test the cat's reflection for consistency by drawing lines between features on the cat and the established reflection vanishing point, and then determining if these lines plausibly pass through reflections of those features. Lines are shown in red for the cat's back, the red strip behind its right ear, the tip of its left ear, and the car it is holding between its forepaws. In each case the lines miss reflected features by a significant margin indicating that the reflection vanishing point consistent with the rest of the scene is inconsistent with the cat and its reflection.

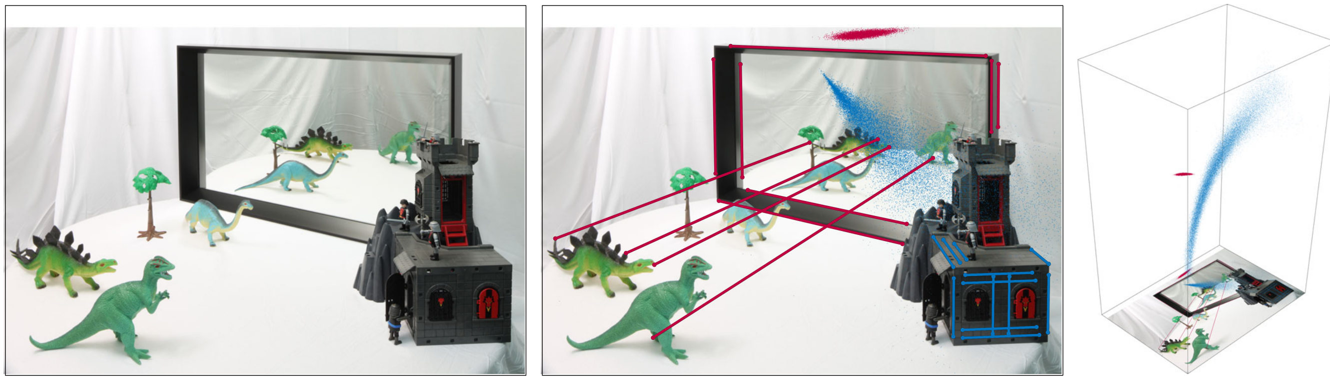

Fig. 10: *The image shown in this figure has been modified by scaling and repositioning the castle. Lines in the center image show features that were used for computing the vanishing points. The resulting clouds of possible locations for the image's center of projection generated from the mirror and from the castle do not overlap and thus reveal the image to be fake.*

Fig. 11: *The first image in the top row shows two objects being reflected separately in two different mirrors. The second and third images in the first row show that the clouds for both the principle points and the centers of projection overlap, indicating that the reflections are mutually consistent. In the bottom row, the left side of the image is from a photograph taken with a focal length different from the right side of the image. The clouds for the principle points still overlap and do not falsify the image. However the clouds for the centers of projection do not overlap, indicating that the image has been manipulated.*

A final example, taken from GQ Magazine, appears in Figure 13. Several features are clearly reflected in the mirror and provide a well-defined reflection vanishing point that is consistent features on the perpendicular wall. While the image appears to be generally legitimate, two inconsistencies reveal the presence of some photo manipulation. First, Reggie Bush does not appear in the mirror and this absence could conceivably be explained by asserting that his reflection is being obscured by Kim Kardashian who is standing in front of the mirror. Indeed, a line from the back of Bush's head to the reflection vanishing point passes through her head and supports this notion. However, we know the reflection line midpoint must fall on the mirror's surface and once that point has been fixed we can determine where his head should appear in the mirror. This location is shown with a blue circle and the absence of his head indicates that the image has been edited. The second inconsistency appears when a line from the top of Kardashian's head is connected

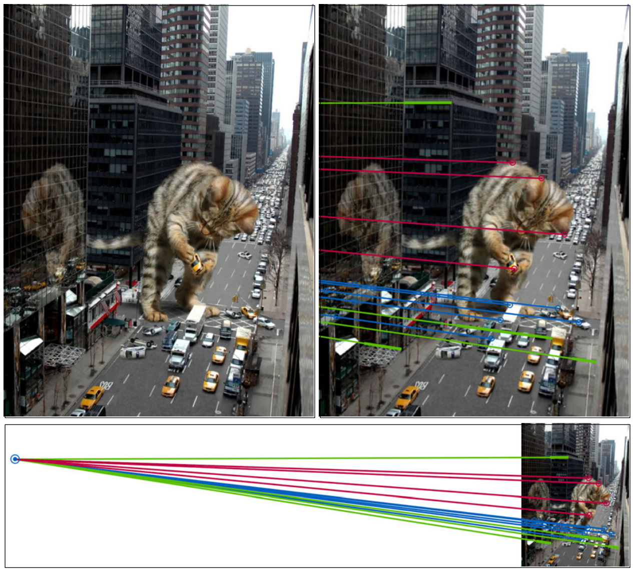

Composite photo World News, copyright 2006.

Fig. 12: *The image in this figure appeared in news articles on April 1 of 2006 and depicts an improbable scenario. Blue lines connect scene features with their corresponding reflections in the building windows. Green lines highlight linear features that should be perpendicular to the building front: an awning, a rooftop, the crosswalk, and a joint in the sidewalk. Extensions of the blue and green lines yields a single, consistent vanishing point. Attempting to connect features on the cat to this same vanishing point using red lines results in those red lines missing the corresponding reflected features by a large margin indicating that the cat is not consistent with the rest of the scene.*

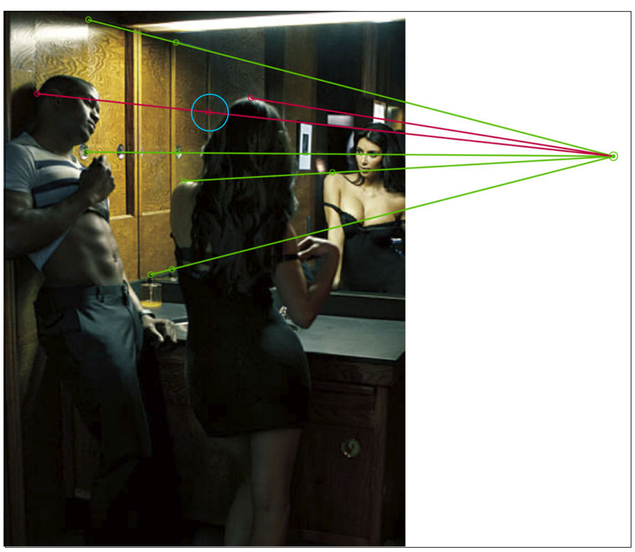

Photo by Alexi Lubomirski, "The Saint and the Sinner," copyright 2009.

Fig. 13: *This image of Reggie Bush and Kim Kardashian appeared in the March 2009 issue of GQ Magazine. The green lines establish a reflection vanishing point for the mirror using several objects in the scene. The red lines reveal that Reggie Bush's head has been removed from the reflection (it should appear in the blue circle) and that the height of Kim Kardashian's hair has been altered.*

to the reflection vanishing point. This line passes above the top of her head in the reflection, indicating that some manipulation has occurred.

5. DISCUSSION

The criteria we have described for testing the authenticity of photographs are appropriate for images containing planar reflectors where parts of the scene can be observed both directly and indirectly through reflection. Although these requirements do limit the applicability of our methods to only certain images, this limitation is expected of any forensic analysis tool. In the cases where our methods do apply, they can provide clear and objective evidence of manipulation. Due to their geometric nature, they are also insensitive to common operations that do not alter scene content such as resampling, color manipulations, and lossy compression.

Mathematica code implementing the methods described here is available for use by others under the terms of a BSD open source license. It can be downloaded from the authors' website.[3]

In comparing centers of projection we explicitly considered the uncertainty inherent in trying to select the exact location of a given feature. This uncertainty can account for both human error as well as blurred or partially occluded features. We did not use this same approach when testing for the existence of a well-defined reflection vanishing point because we did not find it to be necessary. However, the same approach could be used in cases where images are highly blurry or distorted, or where there is reason to suspect human error in feature location. For the cat image we dealt with the difficulty of finding clear features on the cat by simply establishing that no locations on the cat's reflection were consistent with the reflection vanishing point established by the rest of the scene. A cloud-based comparison might be useful in situations were all objects in a scene have poorly-locatable features.

One potential area of future work is to develop similar methods for use with curved surface reflectors. Although very infrequent in natural scenes, man-made environments often contain smooth, curved, and highly specular surfaces such as car bodies, cooking pots, windshields, and plastic appliances. One obvious approach would be to generalize the reflection vanishing point to a curve or surface patch of points. Unfortunately, our initial experimentation with this idea indicates that useful application would require an unreasonable number of corresponding points. Perhaps this limitation can be overcome or some other approach developed.

As with any testable forensic criteria, an informed forger could attempt to fool the test. Indeed, one could imagine a tool based on the analysis in this paper that facilitates creating fake reflections. However based on our experience, it is not trivial to create fake reflections that both look real and that satisfy our criteria. Further, in some instances it is not feasible to create a reflection that would look correct, support the desired fiction of the image, and still pass the tests we have described. For example, in Figure 1 placing the closer figure's reflection to make it consistent with the reflection vanishing point established by the rest of the scene would either make the reflection appear unreasonably large or place the other figure too far away to be making the hand off. Even when it is feasible to create a passable reflection, these tests still create another potential stumbling block that raises the difficulty of creating an undetectable forgery.

[3]http://graphics.berkeley.edu/papers/Obrien-EPM-2012-01

Acknowledgments

The authors thank Ahna Girshick and Marty Banks for their helpful insights and comments. We also thank Torin and Cayden for allowing us the use of their toys.

REFERENCES

ADELSON, E. H. 2000. *The New Cognitive Neurosciences, 2nd Edition*. MIT Press, Chapter Lightness Perception and Lightness Illusions, 339–351.

AVIDAN, S. AND SHAMIR, A. 2007. Seam carving for content-aware image resizing. *ACM Transactions on Graphics (SIGGRAPH) 26,* 3.

BERTAMINI, M., LATTO, R., AND SPOONER, A. 2003. The venus effect: people's understanding of mirror reflections in paintings. *Perception 32*, 593–599.

BERTAMINI, M. AND PARKS, T. E. 2005. On what people know about images on mirrors. *Cognition 98*, 85–104.

BRAVO, M. AND FARID, H. 2001. Texture perception on folded surfaces. *Perception 30,* 7, 819–832.

CAPRILE, B. AND TORRE, V. 1990. Using vanishing points for camera calibration. *Int. J. Comput. Vision 4*, 127–140.

CAVANAGH, P., CHAO, J., AND WANG, D. 2008. Reflections in art. *Spatial Vision 21,* 3-5, 261–270.

CROUCHER, C. J., BERTAMINI, M., AND HECHT, H. 2002. Naive optics: Understanding the geometry of mirror reflections. *Journal of Experimental Psychlogy: Human Perception and Performance 28,* 3, 546–562.

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

FARID, H. AND BRAVO, M. 2010. Image forensic analyses that elude the human visual system. In *SPIE Symposium on Electronic Imaging*. San Jose, CA.

FRIDRICH, J. 2009. Digital image forensic using sensor noise. *IEEE Signal Processing Magazine 26,* 2, 26–37.

FRIDRICH, J., SOUKAL, D., AND LUKAS, J. 2003. Detection of copy move forgery in digital images. In *Digital Forensic Research Workshop*.

GARRY, M. AND GERRIE, M. 2005. When photographs create false memories. *Current Directions in Psychological Science 14*, 326–330.

GARRY, M. AND WADE, K. 2005. Actually, a picture is worth less than 45 words: Narratives produce more false memories than photographs. *Psychonomic Bulletin and Review 12*, 359–366.

GLOE, T., WINKLER, A., AND BOROWKA, K. 2010. Efficient estimation and large-scale evaluation of lateral chromatic aberration for digital image forensics. In *SPIE Conference on Media Forensics and Security*.

GRABLER, F., AGRAWALA, M., LI, W., DONTCHEVA, M., AND IGARASHI, T. 2009. Generating photo manipulation tutorials by demonstration. *ACM Transactions on Graphics (SIGGRAPH) 28,* 3.

HARTLEY, R. AND ZISSERMAN, A. 2004. *Multiple View Geometry in Computer Vision*. Cambridge University Press.

HAYS, J. AND EFROS, A. A. 2007. Scene completion using millions of photographs. *ACM Transactions on Graphics (SIGGRAPH) 26,* 3.

JOHNSON, M. AND FARID, H. 2007a. Detecting photographic composites of people. In *6th International Workshop on Digital Watermarking*. Guangzhou, China.

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. 2006. Exposing digital forgeries through chromatic aberration. In *ACM Multimedia and Security Workshop*. 48–55.

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

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

KELBY, S. 2008. *The Digital Photography Book*. Peachpit Press.

KIRCHNER, M. 2010. Efficient estimation of CFA pattern configuration in digital camera images. In *SPIE Conference on Media Forensics and Security*.

KIRCHNER, M. AND GLOE, T. 2009. On resampling detection in recompressed images. In *IEEE Workshop on Information Forensics and Security*. 21–25.

KUBOVY, M. 1986. *The Psychology of Perspective and Renaissance Art*. Cambridge Univ. Press.

LIGHT, K. 2004. Fonda, Kerry and photo fakery. In *The Washington Post*. February 28, A21.

LUKAS, J., FRIDRICH, J., AND GOLJAN, M. 2006. Digital camera identification from sensor noise. *IEEE Transactions on Information Forensics and Security 1,* 2, 205–214.

MONTAGUE, J. 2010. *Basic Perspective Drawing*, Fifth ed. John Wiley and Sons.

NG, T.-T. AND CHANG, S.-F. 2004. A model for image splicing. In *IEEE International Conference on Image Processing*.

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

PAN, X. AND LYU, S. 2010. Detecting image region duplication using SIFT features. In *International Conference on Acoustics, Speech, and Signal Processing*. 1706–1709.

PETERS, J. W. 2010. On The Economist's cover, only a part of the picture. In *The New York Times*. July 5.

POPESCU, A. C. AND FARID, H. 2004. Exposing digital forgeries by detecting duplicated image regions. Tech. Rep. TR2004-515, Department of Computer Science, Dartmouth College.

POPESCU, A. C. AND FARID, H. 2005a. Exposing digital forgeries by detecting traces of re-sampling. *IEEE Transactions on Signal Processing 53,* 2, 758–767.

POPESCU, A. C. AND FARID, H. 2005b. Exposing digital forgeries in color filter array interpolated images. *IEEE Transactions on Signal Processing 53,* 10, 3948–3959.

POTTER, M. 1976. Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory 2*, 509–522.

RIESS, C. AND ANGELOPOULOU, E. 2010. Scene illumination as an indicator of image manipulation. In *International Workshop on Information Hiding*.

RITSCHEL, T., OKABE, M., THORMÄHLEN, T., AND SEIDEL, H.-P. 2009. Interactive reflection editing. In *ACM SIGGRAPH Asia 2009*. SIGGRAPH Asia '09. 129:1–129:7.

SACCHI, D., AGNOLI, F., AND LOFTUS, E. 2007. Doctored photos and memory for public events. *Applied Cognitive Psychology 21*, 1005–1022.

SHELBOURNE, T. 2007. *Photoshop CS3 Photo Effects Cookbook*. O'Rilly.

SINHA, P., BALAS, B., OSTROVSKY, Y., AND RUSSELL, R. 2006. Face recognition by humans: 19 results all computer vision researchers should know about. *Proceedings of the IEEE 94,* 11, 1948–1962.

SUNKAVALLI, K., JOHNSON, M. K., MATUSIK, W., AND PFISTER, H. 2010. Multi-scale image harmonization. *ACM Transactions on Graphics (SIGGRAPH) 29,* 4.

VISHWANATH, D., GIRSHICK, A., AND BANKS, M. 2005. Why pictures look right when viewed from the wrong place. *Nature Neuroscience 10,* 8, 1401–1410.

WADE, K., GARRY, M., READ, J., AND LINDSAY, D. 2002. A picture is worth a thousand lies. *Psychonomic Bulletin and Review 9*, 597–603.

WADE, N. 2005. Clone scientist relied on peers and Korean pride. In *The New York Times*. December 25.

WESTHEIMER, G. AND MCKEE, S. 1975. Visual acuity in the presence of retinal-image motion. *Journal of the Optical Society of America 65*, 847–850.

APPENDIX

A. COMPUTING THE CENTER OF PROJECTION

The center of projection, $\mathbf{C}$, for an image where three mutually orthogonal vanishing points, $\mathbf{V}_i$, have been identified must satisfy the following system of quadratic equations:

$$(\mathbf{C}-\mathbf{V}_1)\cdot(\mathbf{C}-\mathbf{V}_2)=0 \tag{10}$$
$$(\mathbf{C}-\mathbf{V}_2)\cdot(\mathbf{C}-\mathbf{V}_3)=0 \tag{11}$$
$$(\mathbf{C}-\mathbf{V}_3)\cdot(\mathbf{C}-\mathbf{V}_1)=0 \quad . \tag{12}$$

Although this system of equations is quadratic in $\mathbf{C}$, it can be solved easily by reducing to a linear system of equations.

Each equation is expanded to obtain:

$$\mathbf{C}\cdot\mathbf{C}-\mathbf{C}\cdot(\mathbf{V}_1+\mathbf{V}_2)+\mathbf{V}_1\cdot\mathbf{V}_2=0 \tag{13}$$
$$\mathbf{C}\cdot\mathbf{C}-\mathbf{C}\cdot(\mathbf{V}_2+\mathbf{V}_3)+\mathbf{V}_2\cdot\mathbf{V}_3=0 \tag{14}$$
$$\mathbf{C}\cdot\mathbf{C}-\mathbf{C}\cdot(\mathbf{V}_3+\mathbf{V}_1)+\mathbf{V}_3\cdot\mathbf{V}_1=0 \quad . \tag{15}$$

The only quadratic term, $\mathbf{C}\cdot\mathbf{C}$ is common to all three equations and can be eliminated by defining an auxiliary variable, $d=\mathbf{C}\cdot\mathbf{C}$. We also select a coordinate system such that the plane containing the $\mathbf{V}_i$ is the $Z=0$ plane. Doing so eliminates C_z from the dot-products with the $\mathbf{V}_i$ yielding:

$$d-C_x(V_{1x}+V_{2x})-C_y(V_{1y}+V_{2y})+\mathbf{V}_1\cdot\mathbf{V}_2=0 \tag{16}$$
$$d-C_x(V_{2x}+V_{3x})-C_y(V_{2y}+V_{3y})+\mathbf{V}_2\cdot\mathbf{V}_3=0 \tag{17}$$
$$d-C_x(V_{3x}+V_{1x})-C_y(V_{3y}+V_{1y})+\mathbf{V}_3\cdot\mathbf{V}_1=0 \tag{18}$$

which is linear in the three unknowns, d, C_x, and C_y. Once they have been solved for, C_z can be recovered using

$$C_z=\sqrt{d-C_x^2+C_y^2} \quad . \tag{19}$$

The sign ambiguity corresponds to a pair of solutions on either side of the image plane. An imaginary solution indicates that the three spheres do not all intersect and that the three vanishing points are inconsistent with mutually orthogonal lines in a real image.