Illumination Analysis in Physics-based Image Forensics: A Joint Discussion of Illumination Direction and Color

Christian Riess

IT Security Lab,
University of Erlangen-Nuremberg, Germany

Abstract. Illumination direction and color are two physics-based forensic cues that are based on the same underlying model. In this work, we discuss these methods in the light of their joint physical model, with a particular focus on the limitations and a qualitative study of failure cases of these methods. Our goal is to provide directions for future research to further reduce the list of constraints that these methods require in order to work. We hope that this eventually broadens the applicability of physics-based methods, and to spread their main advantage, namely their stringent models for deviations of the expected image formation.

Keywords: Blind passive image forensics, physics-based image forensics, illumination environments, illumination color

1 Introduction

The goal of image forensics is to provide algorithms for determining authenticity and origin of image data [12, 24, 27]. With the increasing ease of creating and distributing digital imagery, attention to this field of research also increased in recent years. Within image forensics, our particular interest for this paper is in the task of detecting image manipulations. The majority of the proposed approaches to manipulation detection arguably belongs to the group of so-called statistical methods. These algorithms look for statistical deviations from an expected model that is built around e. g.,, compression artifacts [1, 3], noise residuals [2, 4, 7], resampling artifacts [19, 23], or copy-move forgeries [5, 6]. Advantages of the statistical methods are that they can typically run offline in batch processing large amounts of images, and that they achieve excellent detection rates, particularly in scenarios where the amount of postprocessing to the imagery is limited.

Another, less popular group of methods in image forensics are the so-called physics-based approaches. Here, a model is built on a physical property of the scene [12]. Deviations from the model are assumed to be manipulations. Notable features to physics-based methods are illumination direction [15, 17, 22], illumination color [9, 13, 25], specular reflection [16, 20, 31], vanishing points [14], and shadows [18, 32]. Broadly speaking, physics-based methods are widely accepted as being less sensitive to image postprocessing like compression or downsampling. Also, the fact that the underlying models are typically drawn from physics or optics, allows in principle stronger statements on the (in-)consistency: a physics-based method typically constrains the observed interaction between scene elements or parts of scene elements. However, as a consequence, the applicability and chances for success of a physics-based method very often depends on the actual content of the scene. This dependence on scene content is oftentimes perceived as the most serious limitation of physics-based methods. The second severe limitation of many physics-based approaches is that they require user interaction to reliably work. Thus, most of the current methods can not easily be adapted to an automated batch processing pipeline.

Algorithms that exploit the interaction of illumination with the scene take a central position among physics-based methods. Our group has investigated in recent years manipulation cues from investigating illumination direction and illumination color. In this work, we provide an overview of our findings and our perspective on open issues in illumination-based forensics. By "illumination-based forensics", we denote forensic methods that operate on objects under direct illumination (as opposed to, for example, methods that explicitly make use of shadow boundaries [18]). We first introduce a general model for image formation in Sec. 2.1, and discuss this model with respect to illumination-based forensics in Sec. 2.2. The approaches to exploiting illumination environments (or direction, respectively) are presented in Sec. 3, including some notable limitations. A similar presentation is done in Sec. 4 for the forensic exploitation of illumination color. A short discussion and conclusion on the similarities and differences between both approaches is presented in Sec. 5

2 Illumination-based Forensics

We first introduce the overall model of image formation in Sec. 2.1, and then discuss its application to image forensics in Sec. 2.2.

2.1 Image Formation

The general image formation model for a color camera is

$$\boldsymbol{i}(\boldsymbol{x}, \boldsymbol{\theta}) = \int\limits_{\Omega} \int\limits_{\lambda} r(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta}) \cdot e(\boldsymbol{x}, \boldsymbol{\nu}, \lambda) \cdot c(\lambda) \, \mathrm{d}\lambda \, \mathrm{d}\boldsymbol{\nu} \ , \tag{1}$$

where $\boldsymbol{x}$ denotes a point in the scene and $\boldsymbol{\theta}$ subsumes geometric scene parameters, most notably the camera pose. $e(\boldsymbol{x}, \boldsymbol{\nu}, \lambda)$ denotes the intensity of incident light on $\boldsymbol{x}$ from direction $\boldsymbol{\nu}$ with wavelength λ. $r(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta})$ denotes the reflectance function. Finally, $\boldsymbol{c}(\lambda)$ denotes the vector of camera responses for light of wavelength λ, which essentially models the red, green, and blue color filters of the camera. In the case of a grayscale camera, $\boldsymbol{c}(\lambda)$ is just a one-dimensional vector containing a single broad color filter. The relationship of these quantities is also sketched in Fig. 1. Here, the incident spectrum of wavelengths λ is shown in magenta. The spectrum potentially varies across the hemisphere of incident angles

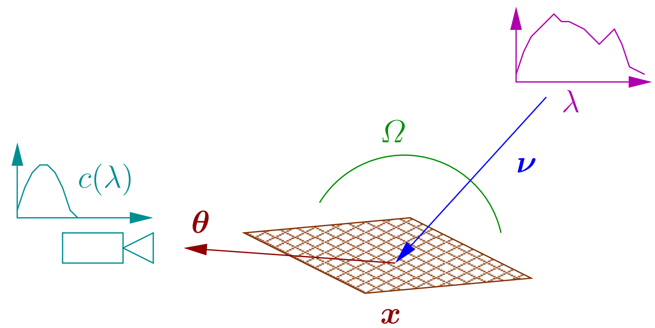

Fig. 1. Sketch for the set of angles of incident light Ω with individual directions $\boldsymbol{\nu}$, the angle of exiting light $\boldsymbol{\theta}$, and the wavelength of light λ incident on a surface location $\boldsymbol{x}$, analogously to the image formation model of Eqn. 1.

Ω (green), where $\boldsymbol{\nu}$ (blue) denotes one specific direction of incidence of light onto a surface location $\boldsymbol{x}$. The reflectance function at $\boldsymbol{x}$ is evaluation with respect to the direction of exitance $\boldsymbol{\theta}$ (red). The contribution of each color to the respective color channel is determined by the color filters $c(\lambda)$ (cyan). Note that this model only covers photometric effects, but does not include any non-linearities. For image forensics, the probably most important non-linear contributions are the camera's gamma function and JPEG compression. We omitted these effects from Eqn. 1 because these effects are introduced by the digital signal processor, i.e., after the physical image formation.

The image formation model in Eqn. 1 is rather general. Specializations can be done by specifying the reflectance function $r(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta})$. For example, the widely known Lambertian reflectance model can be expressed as

$$r_{\text{Lambert}}(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta}) = \rho(\boldsymbol{x}, \lambda) \cdot \max\left(\frac{\boldsymbol{n}(\boldsymbol{x}) \circ \boldsymbol{\nu}}{\|\boldsymbol{n}(\boldsymbol{x})\|\|\boldsymbol{\nu}\|}, 0\right) \quad , \tag{2}$$

where $\rho(\boldsymbol{x}, \lambda)$ denotes the albedo for wavelength λ at position $\boldsymbol{x}$, and $\boldsymbol{n}(\boldsymbol{x})$ denotes the surface normal at point $\boldsymbol{x}$ and $\circ$ denotes the scalar product. This reflectance model is also referred to as "purely diffuse" reflectance, since the reflected intensity only depends on the angle between incident light and surface normal, but not on the position of the viewer. Another alternative is the dichromatic reflectance model by Shafer [28]. Here, the reflectance term splits into a diffuse and a specular portion,

$$r_{\text{Dichrom}}(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta}) = \alpha r_{\text{Lambert}}(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta}) + (1 - \alpha) r_{\text{spec}}(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta}) \quad , \tag{3}$$

where the weighting term α depends on the surface material at $\boldsymbol{x}$. The characteristic of the specular component r_{spec} is that it is zero everywhere except of the specific constellation where angle of incidence $\boldsymbol{\nu}$, surface normal $\boldsymbol{n}(\boldsymbol{x})$ and the vector of the viewing direction $\boldsymbol{\theta}$ are coplanar, and the angles $(\boldsymbol{\nu}, \boldsymbol{n}(\boldsymbol{x}))$ and $(\boldsymbol{n}(\boldsymbol{x}, \boldsymbol{\theta})$ are identical. For color research, one additional, common assumption for the dichromatic reflectance model is the so-called neutral interface reflectance assumption. It states that the observed color in a pure specularity is identical to the color of the light source.

2.2 Forensic Applications of the Image Formation Model

The key assumption of illumination-based forensics is that with an appropriate choice of the reflectance function $r(\boldsymbol{x}, \boldsymbol{\nu}, \lambda, \boldsymbol{\theta})$, and sufficient information about the environment (e.g., about number and positions of light sources, or object surface normals), it is possible to explain the observed pixels with the model. In particular, this also allows to compare two objects with respect to the quality of the model fit. Thus, if an image is manipulated, e. g.,, by splicing it from two sources, it may be the case that one object matches the model while the other one does not. Such a disagreement between objects within the same image is then taken as an indicator for manipulation.

In practice, existing methods estimate the model parameters from each object (as if the object would be an original, and belong to that scene). In a second step, model parameters between objects are compared, and outliers are marked as manipulated. In the context of Eqn. 1, any of the explicitly or implicitly given model parameters can be used for detecting inconsistencies. For example, methods that are based on the illumination direction estimate the surface normal $\boldsymbol{n}(\boldsymbol{x})$ to compute the intensity distribution of the lighting environment $e(\boldsymbol{x}, \boldsymbol{\nu}, \lambda)$, integrated over all λ in the visible range and all $\boldsymbol{x}$ in a sufficiently local neighborhood. Methods based on the illumination color are in a less apparent way linked to Eqn. 1, but they basically exploit small variations in the surface normals $\boldsymbol{n}(\boldsymbol{x})$ to estimate $e(\boldsymbol{x}, \boldsymbol{\nu}, \lambda) \cdot c(\lambda)$, integrated over all $\boldsymbol{\nu}$, and all $\boldsymbol{x}$ in a sufficiently local neighborhood.

3 Illumination Direction

The key ideas for examining illumination direction were presented by Johnson and Farid [15] and Kee and Farid [17]. Both works propose ways to estimate the distribution of incident light $e(\boldsymbol{x}, \boldsymbol{\nu}, \lambda)$ using estimates of the surface normals $\boldsymbol{n}(\boldsymbol{x})$. At object contours, the z-component of surface normals is 0, i. e.,, the normals are located in the image plane. Johnson and Farid propose to estimate the orientation of these normals from orientation of the contour. Assuming furthermore Lambertian reflectance on a monochromatic image, and user-selected directly illuminated (non-shadowed) contours along the same surface material, the estimation of $e(\boldsymbol{x}, \boldsymbol{\nu}, \lambda)$ can be written as a least squares fit of the observed contour intensities to a spherical harmonics model [15].

Specifically, the estimation problem can be written as

$$A\boldsymbol{h} = \boldsymbol{i} \ , \tag{4}$$

where $\boldsymbol{i} = (i(\boldsymbol{x}_1), \ldots, i(\boldsymbol{x}_N))^{\mathrm{T}}$ denotes the N-element vector of observed (gray-level) intensities along the contour, and $A \in \mathbb{R}^{n \times 5}$ denotes a matrix of spherical harmonics bases that are evaluated for the surface normals $(n(\boldsymbol{x}_1), \dots, n(\boldsymbol{x}_N))$ associated with each of the N contour points. The unknown vector $\boldsymbol{h} \in \mathbb{R}^{5\times 1}$ consists of the five coefficients that parameterize a second-order 2-D spherical harmonics model. For increased robustness, this linear system of equations is solved using a Tikhonov regularizer [15].

The same idea is used in the work by Kee and Farid [17], with the main exception that an *a priori* known object class — faces — is used to estimate the surface normals. Due to major progress in recent years in fitting a 3-D face model to a 2-D face image, it became possible to obtain 3-D surface normals for all face pixels instead of only the 2-D contour normals. This allows to estimate 3-D lighting environments, with essentially the same fitting approach as above. The main difference is that a second-order 3-D model consists of nine coefficients instead of five.

There are two immediate advantages from the availability of 3-D normals. First, since all face pixels can be used instead of only a few contour pixels, the numerical estimation of the lighting environment becomes much more robust. Second, the 3-D model allows to not only distinguish 2-D projections of light sources onto the image plane, but their location in 3-D space. Both factors together make the 3-D approach more accurate than the 2-D approach, at the expense of being only applicable to scenes with at least two faces.

3.1 Addressing the Constraints: Surface Albedo, Automation and Shape-from-Shading

With the outlined basic idea published, a number of follow-up works aim at reducing the constraints that are imposed on the methods. For example, for numerical stability of the 2-D contour-based approach, it is necessary to collect occluding contours from an angular range of about 120 degrees or more. The brightness information along the contours must be obtained from materials with the same object color (albedo). In our work, we found that this requirement of constant albedo is one of the big obstacles for using the method in practice. When selecting in an image contours along two objects for comparison, we regularly were unable to find admissible, monochromatic occluding contours across the required angular range. Thus, we proposed to include a multiplicative factor to the estimated coefficients that neutralizes surface colors [26]. A similar idea has been proposed by Peng et al. for the 3-D approach, namely to obtain a texture estimate together with the face fit that is used to normalize the intensity distribution [22]. Furthermore, since face detection is a well-examined computer vision task, the same group also proposed a slight extension of the original method that automatically finds faces and fits a face model to it [21]. Another severe restriction of the 3-D method is that it can only detect manipulations that involve faces. This has been addressed by Fan et al. by experimenting with shape from shading on more general objects [11]. Nevertheless, this additional flexibility naturally comes at the expense of a somewhat reduced model accuracy, which, by extension, reduces the ability of the method to detect manipulations.

Fig. 2. Examples for different surface materials and textures. Top left: synthetic materials like outdoor clothing typically violate the Lambertian reflectance model. Top right: unstructured materials like hair are very prone to mislead the estimation, and hence need to be segmented out beforehand. Bottom row: texture like on the jeans or the makeup is oftentimes barely separable from intensity gradients that arise from varying illumination.

3.2 Unsolved Challenges in Lighting Environments for Forensics

We noticed in our experiments with 2-D and 3-D estimators for lighting environments, that there also exists a number of "hidden" constraints. These constraints can be readily derived from the underlying model, but we find them to be somewhat less obvious. We group these constraints into surface materials/texture, geometry, and scene layout.

Surface Material and Texture: Many artificial materials exhibit unusual reflectance behavior, that greatly differs from the assumed Lambertian reflectance model. An example case is shown in Fig. 2 in the top left. While casual clothing is oftentimes approximated as Lambertian, this does not hold for synthetic outdoor clothing like the shown ski suit. Also, hair (Fig. 2, top right) and smooth textures on the surface (Fig. 2, bottom row) are cases where material and illumination information are easily confused, leading to failure cases. An approach to mitigate these issues in an automated manner might be to use specialized detectors for the most common cases. For example, the color gradients on the jeans in Fig. 2 on the bottom left is actually a colored gradient, and could probably as such be filtered out.

Fig. 3. Examples for object geometries that do not permit a reliable 3-D estimation of the direction of the incident illuminant. The reason is that these objects do not exhibit normals in all directions, but only along a lower-dimensional manifold. In the left and middle image, the columns and the hangar roof spread out the surface normals along a 2-D plane. In the right image, the flat building facade confines all surface normals to be parallel. As a result, the observations are severely undersampled and can not be used for a stand-alone estimation of the illuminant direction.

Geometry: Scene geometry can cause subtle, yet significant errors in the estimation of the illuminant direction. One requirement of the method is that a sufficient variety of surface normals is observed. However, many man-made structures exhibit surfaces that reside on a lower-dimensional manifold, and hence can not provide a sufficient variety of surface normals. Example images are shown in Fig. 3. On the left, the surface normals on the columns all reside on a plane that perpendicularly intersects the columns. As a consequence, a 3-D lighting estimator that solely relies on these normals is unable to distinguish different elevation angles of the illuminant. An almost identical examples is shown in the middle image for the cylindrical hangar roof. A third example, where all surface normals are even parallel, is the building facade shown in the right image. It is important to note, and to communicate to users of lighting-environment methods that in such cases the scene geometry inherently limits the applicability of the method, even if the normals on these problematic surfaces are admissible with respect to the other required constraints.

Scene Layout: Upon closer examination, it turns out that a surprising number of scene layouts, i. e.,, constellations of objects, cameras, and light sources, create very challenging cases for the use of methods based on the direction of the illumination. The canonical best case for the analysis of 2-D lighting environments are light sources that are located between camera and object. Since the 2-D approach is only able to distinguish the projection of the lighting environment onto the image plane, analysis becomes in principle easier if the light sources are located closer to the object than to the camera. A good case in terms of lighting positioning for the 2-D estimator is shown in Fig. 4 in the top left. Three other typical cases are shown in the remaining pictures from Fig. 4. In the top right, the light is mainly coming from the front. This is a case where the 2-D estimator becomes less accurate. However, the 3-D estimator can operate on images of this type. In the bottom left, the light source is located behind the person, which is not modelled by the methods and hence is a situation where the method can not

Fig. 4. Examples for scene geometry and lighting. Top left: with a single dominant light source between camera and object, in a lateral position, this is a case where the 2-D lighting estimator is expected to perform very reliably. Top right: light sources located more in front of the object are challenging for the 2-D estimator, but can be estimated with the 3-D lighting estimator. Bottom left: the light source is located behind the object, which is modelled by neither of the two estimators. Bottom right: another failure case for both estimators is the absence of a direct light source, e. g.,, in a shadow region.

be used. Similarly in the bottom right, where there is apparently no directional light, but only ambient illumination.

Not shown in this image, but in Fig. 6, are two cases where the scene consists of more than a single light source. This is no problem if the mixture of light sources is identical on the compared objects or persons. However, the problem becomes ill-posed in the case of multiple inhomogeneous light sources, which makes the two lighting environments incomparable [22].

4 Illumination Color

A second approach to use Eqn. 1 for the detection image manipulations is to solve for the spectral distribution of the light source $e(\boldsymbol{x}, \boldsymbol{\nu}, \lambda) \cdot c(\lambda)$. Gholap and Bora proposed to examine objects that can be assumed to follow the dichromatic reflectance model, and to use the intersection point of so-called “dichromatic lines” to detect image manipulations [13]: if each of the three objects is illuminated by the same light source, the three dichromatic lines are expected to intersect in one point.

In our work, we followed a slightly different approach. We propose to subdivide the image into small regions, and to estimate the color of the illuminant locally for each region [25]. In computer vision, estimating the color of the illuminant is typically a preparatory step to white balancing (the field of research is called "color constancy"). However, in a forensic context, we assume that the image is manipulated if these local estimates differ too much. The color constancy estimators with the lowest error rates are statistical, and many of these are based on machine learning. Such estimators are not optimal from a theoretical perspective: if a statistical estimator is applied locally to an image region, the area may just not contain sufficiently diverse pixel information in order to allow a low-error estimate. Second, the statistical estimators typically lack a theoretical model for their success cases and failure cases. In that sense, an illuminant color framework containing statistical estimators is not entirely physics-based anymore. To mitigate these issues, we preferably work with physics-based illuminant color estimator by Tan et al. that assumes the dichromatic reflectance model [29]. This estimator requires surface patches that exhibit some variation in the surface normals. However, the requirements are much milder than for the set of surface normals that are needed for estimating an illumination environment. However, the method additionally requires that some of the pixels exhibit partially specular reflectance. It turns out that this requirement is a major source of error in many practical scenarios, particularly for pixel regions in the background. However, in our previous work, we paired the physics-based estimator with a statistical estimator, with the goal to backup the physics-based estimator in such cases [9].

An interesting idea for the purpose of source camera identification has been proposed by Deng et al., by assuming that repeated white balancing with the camera's internal white balancing algorithm is an idempotent operation [10]. This allows to find a camera-model specific white balancing algorithm for a given input image.

4.1 Addressing the Constraints: Surface Albedo and Learned Decision Boundaries

In our experiments, it turned out that a direct, agnostic comparison of locally estimated illuminant colors is oftentimes a too weak decision criterion. The estimation of an illuminant color is an underconstrained problem, and as such must rely on constraints to be solvable. Estimations on local image patches, as proposed for the purpose of image forensics, by tendency further increases the error rates. We observed that particularly the object color influences the direction of the illuminant estimate. As a consequence, we propose a method that operates on areas where we can assume to observe roughly the same underlying material, namely on faces [9]. The idea is to relate illuminant estimates from similar underlying objects to be able to set up a tighter decision boundary. In this work, the decision is learned from a rather sophisticated feature map, which unfortunately leaves the physics-based grounds of the method. Nevertheless, this closes the gap towards a more reliable detection method based on illuminant color.

Fig. 5. Example face shots exhibiting small specularities. Physics-based illuminant estimators oftentimes benefit from specular areas. Instead of estimating these regions directly, it may be beneficial to heuristically search for objects like faces, expecting that nose, cheeks or forehead exhibit slight specularities.

4.2 Unsolved Challenges in Color Constancy for Forensics

Although the underlying ideas for using illuminant color in image forensics have already been presented a while ago, there are — particularly on the theoretical side — a number of unsolved questions with respect to the realization of an end-to-end physics-based pipeline. Thus, the purpose of this section is to provide a brief introduction into what we belief are the major open issues that need to be addressed.

Dependence on Specularities: Most physics-based illuminant color estimators expect at least partially specular pixels to perform their estimation. Conversely, if a region contains only purely diffuse pixels, it is most likely that a specularity-based estimator will fail. However, specularity segmentation is in general similarly hard as the color constancy problem itself. Nevertheless, it may be an option to use prior information from another channel to constrain the domain of where specularities are more likely to be found. For example, oftentimes exhibit faces a slight specular spots at the tip of the nose. Example pictures are shown in Fig. 5. In such a case, it may suffice to find faces and search within the face for specularities instead of blindly segmenting the whole image for specularities.

Physics-based Decisions: One major missing link for a fully, end-to-end physics-based method that exploits the illuminant color is a physically motivated, yet reliable, metric to distinguish consistent areas from inconsistent areas. Several results have been reported for statistical decision functions that build on top of the illuminant color estimators [8, 9, 30]. Therefore, it might be a good approach to search for general, ideally analytical, rules for distinguishing consistent from inconsistent illuminations.

Scene Layout: Analogously to the analysis of illumination environments, scene layout, and particularly local illuminants have great impact on the success of the illuminant estimators. For example, a considerable number of indoor scenes are

Fig. 6. Examples for multiple illuminants. Indoor scenes with windows are common multi-illuminant cases. Another case (not shown here) are pictures taken with a flash light.

multi-illuminant scenes, with sunlight from a window, another electric light, and maybe even additional secondary light sources, e. g.,, from electronic devices, or interreflections from furniture or walls. Two examples for such cases are shown in Fig. 6. Another common source for two illuminants are photographs that are taken with flash light that also include some background illumination. Resolving cases with multiple different, inhomogeneous illuminants can lead to very difficult situations, since the task of comparing illumination environments is not well-defined anymore: a difference in illuminant colors can either be due to a spliced manipulation, or due to the fact that there are two different light sources in the scene. It turns out that this problem is exactly analogous to multiple light sources for estimating illuminant environments at the end of Sec. 3.2.

5 Discussion and Conclusion

We provide an overview of physics-based methods that analyze illumination environments and illumination color for the purpose of exposing image manipulations. Physics-based methods share the advantage that they are based on rigorous models, but they suffer from their relatively strict assumptions on the content of the scene. This inherently limits their applicability, particularly in an automated batch-processing scenario.

In our review, we show that both families of physics-based methods are based on essentially the same models for image formation, but both families focus on different quantities in the model. For estimating illumination environments, the surface normals of an object are required. For estimating the illuminant color, it is required that an object exhibits a certain variation in the surface normals, but less than for illumination environments. Additionally, for an end-to-end physics-based model, it is required that some pixels exhibit partially specular reflectance to estimate the illuminant color. Thus, both families of methods have similar, yet somewhat different requirements on the object areas they operate on.

As a consequence, estimators for illumination environments have stricter requirements on a sufficiently diverse geometric structure of the object. For example, all normal vectors of a column reside in a plane, and hence the estimated illuminant environment is also confined to that plane. Conversely, estimators for the illuminant color only require a small variation of the surface normal. However, physics-based estimators typically require at least partially specular pixels, which are oftentimes only very sparsely contained in a scene, and may be difficult to segment automatically. We also provide several other pointers to ongoing challenges with both estimators.

Since both methods are based from the same underlying formulation, it may be interesting to further explore the consequences of this joint lineage. For example, it appears to be interesting to search for a confidence measure that is derived from the methods' dependence on the underlying geometry. It may also be interesting to further, more rigorously investigate in which cases each of the two approaches can succeed, and to quantify their overlap in applicability to forensic examination of images.

Acknowledgements

This material is based on research sponsored by the Air Force Research Laboratory and the Defense Advanced Research Projects Agency under agreement number FA8750-16-2-0204. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory and the Defense Advanced Research Projects Agency or the U.S. Government.

References

1. Barni, M., Chen, Z., and Tondi, B. Adversary-aware, data-driven detection of double JPEG compression: How to make counter-forensics harder. In IEEE International Workshop on Information Forensics and Security (Dec. 2016).
2. Bayar, B., and Stamm, M. C. A Deep Learning Approach to Universal Image Manipulation Detection Using a New Convolutional Layer. In ACM Workshop on Information Hiding and Multimedia Security (June 2016), pp. 5–10.
3. Bianchi, T., Piva, A., and Perez-Gonzalez, F. Near Optimal Detection of Quantized Signals and Application to JPEG Forensics. In IEEE International Workshop on Information Forensics and Security (Nov. 2013), pp. 168–173.
4. Chen, M., Fridrich, J., Goljan, M., and Lukás, J. Determining image origin and integrity using sensor noise. IEEE Transactions on Information Forensics and Security 3, 1 (Mar. 2008), 74–90.
5. Christlein, V., Riess, C., Jordan, J., Riess, C., and Angelopoulou, E. An Evaluation of Popular Copy-Move Forgery Detection Approaches. IEEE Transactions on Information Forensics and Security 7, 6 (Dec. 2012), 1841–1854.
6. Cozzolino, D., Poggi, G., and Verdoliva, L. Efficient Dense-Field CopyMove Forgery Detection. IEEE Transactions on Information Forensics and Security 10, 11 (Nov. 2015), 2284–2297.

7. Cozzolino, D., Poggi, G., and Verdoliva, L. Splicebuster: a new blind image splicing detector. In IEEE International Workshop on Information Forensics and Security (2015), pp. 1–6.
8. de Carvalho, T. J., Faria, F. A., Pedrini, H., Torres, R. S., and Rocha, A. Illuminant-based Transformed Spaces for Image Forensics. IEEE Transactions on Information Forensics and Security 11, 4 (Dec. 2015), 720–733.
9. de Carvalho, T. J., Riess, C., Angelopoulou, E., Pedrini, H., and Rocha, A. Exposing Digital Image Forgeries by Illumination Color Classification. IEEE Transactions on Information Forensics and Security 8, 7 (2013), 1182–1194.
10. Deng, Z., Gijsenij, A., and Zhang, J. Source Camera Identification using Auto-White Balance Approximation. In Proceedings of the 13th IEEE International Conference on Computer Vision (ICCV 2011) (Barcelona, Spain, Nov. 2011), pp. 57–64.
11. Fan, W., Wang, K., Cayre, F., and Xiong, Z. 3D Lighting-Based Image Forgery Detection using Shape-from-Shading. In Proceedings of the 20th European Signal Processing Conference (EUSIPCO-2012) (Bucarest, Romania, Aug. 2012), pp. 1777–1781.
12. Farid, H. Photo Forensics. MIT Press, 2016.
13. Gholap, S., and Bora, P. K. Illuminant Colour Based Image Forensics. In IEEE Region 10 Conference TENCON (TENCON 2008) (Hyderabad, India, Nov. 2008).
14. Iuliani, M., Fanfani, M., Colombo, C., and Piva, A. Reliability Assessment of Principal Point Estimates for Forensic Applications. Journal of Visual Communication and Image Representation 42 (1 2017), 65–77.
15. Johnson, M., and Farid, H. Exposing Digital Forgeries in Complex Lighting Environments. IEEE Transactions on Information Forensics and Security 2, 3 (Sept. 2007), 450–461.
16. Johnson, M., and Farid, H. Detecting Photographic Composites of People. In Proceedings of the 7th International Workshop on Digital Watermarking (IWDW 2008) (Busan, South Korea, Dec. 2008), vol. Lecture Notes in Computer Science 5450, pp. 19–33.
17. Kee, E., and Farid, H. Exposing Digital Forgeries from 3-D Lighting Environments. In Proceedings of the 2nd IEEE International Workshop on Information Forensics and Security (WIFS 2010) (Seattle, WA, USA, Dec. 2010).
18. Kee, E., O'Brien, J. F., and Farid, H. Exposing Photo Manipulation from Shading and Shadows. ACM Transactions on Graphics 33, 5 (Aug. 2014), 165:1–21.
19. Kirchner, M. Linear Row and Column Predictors for the Analysis of Resized Images. In ACM SIGMM Multimedia & Security Workshop (Sept. 2010), pp. 13–18.
20. O'Brien, J. F., and Farid, H. Exposing Photo Manipulation with Inconsistent Reflections. ACM Transactions on Graphics 31, 1 (Jan. 2012), 1–11.
21. Peng, B., Wang, W., Dong, J., and Tan, T. Automatic Detection of 3-D Lighting Inconsistencies via a Facial Landmark based Morphable Model. In IEEE International Conference on Image Processing (2016), pp. 3932–3936.
22. Peng, B., Wang, W., Dong, J., and Tan, T. Optimized 3D Lighting Environment Estimation for Image Forgery Detection. IEEE Transactions on Information Forensics and Security 12, 2 (Feb. 2017), 479–494.
23. Popescu, A. C., and Farid, H. Exposing Digital Forgeries by Detecting Traces of Resampling. Signal Processing 53, 2 (Feb. 2005), 758–767.
24. Redi, J., Taktak, W., and Dugelay, J.-L. Digital Image Forensics: A Booklet for Beginners. Multimedia Tools and Applications 51, 1 (Jan. 2011), 133–162.

25. Riess, C., and Angelopoulou, E. Scene Illumination as an Indicator of Image Manipulation. In Proceedings of the 12th International Conference on Information Hiding (IH 2010) (Calgary, AB, Canada, June 2010), vol. Lecture Notes in Computer Science 6387, pp. 66–80.
26. Riess, C., Unberath, M., Naderi, F., Pfaller, S., Stamminger, M., and Angelopoulou, E. Handling Multiple Materials for Exposure of Digital Forgeries using 2-D Lighting Environments. Multimedia Tools and Applications 76, 4 (2016), 4747–4764.
27. Sencar, H. T., and Memon, N., Eds. Digital Image Forensics. Springer, 2013.
28. Shafer, S. A. Using Color to Separate Reflection Components. Journal Color Research and Application 10, 4 (Winter 1985), 210–218.
29. Tan, R., Nishino, K., and Ikeuchi, K. Color Constancy through Inverse-Intensity Chromaticity Space. Journal of the Optical Society of America A 21, 3 (Mar. 2004), 321–334.
30. Wu, X., and Fang, Z. Image Splicing Detection Using Illuminant Color Inconsistency. In Proceedings of the 3rd IEEE International Conference on Multimedia Information Networking and Security (MINES 2011) (Shanghai, China, Nov. 2011), pp. 600–603.
31. Yu, H., Ng, T.-T., and Sun, Q. Recaptured Photo Detection using Specularity Distribution. In Proceedings of the 15th IEEE International Conference on Image Processing (ICIP 2008) (San Diego, CA, USA, Oct. 2008), pp. 3140–3143.
32. Zhang, W., Cao, X., Zhang, J., Zhu, J., and Wang, P. Detecting Photographic Composites using Shadows. In Proceedings of the IEEE International Conference on Multimedia and Expo (ICME 2009) (Cancun, Mexico, June 2009), pp. 1042–1045.