Detection of tampering in the Digital Image Using Illumination Color Estimator

T.Selvi (Assistant Professor), Swetha.T, Narmadha.M
Computer Science Department, Kings Engineering College, Chennai.

Abstract—For decades, photographs have been used to document space-time events and they have often served as evidence in courts. Although photographers are able to create composites of analog pictures, this process is very time consuming. Today, however, powerful digital image editing software makes image modifications straightforward. This undermines our trust in photographs and, in particular, questions pictures as evidence for real-world events. In this paper, we analyze one of the most common forms of photographic manipulation, known as image composition or splicing. We propose a forgery detection method that exploits subtle inconsistencies in the color of the illumination of images. Our approach is machine-learning based and requires minimal user interaction. The technique is applicable to images containing two or more people and requires no expert interaction for the tampering decision. To achieve this, we incorporate information from physics- and statistical-based illuminant estimators on image regions of similar material. From these illuminant estimates, we extract texture- and edge-based features which are then provided to a machine-learning approach for automatic decision-making. The classification performance using an SVM meta-fusion classifier is promising. It yields detection rates of 86% on a new benchmark dataset consisting of 200 images, and 83% on 50 images that were collected from the Internet.

I. Introduction

Every day, millions of digital documents are produced by a variety of devices and distributed by newspapers, magazines, websites and television. In all these information channels, images are a powerful tool for communication. Unfortunately, it is not difficult to use computer graphics and image processing techniques to manipulate images. How we deal with photographic manipulation raises a host of legal and ethical questions that must be addressed [1]. However, before thinking of taking appropriate actions upon a questionable image, one must be able to detect that an image has been altered.

Fig. 1. How can one assure the authenticity of a photograph? Example of a spliced image involving people.

Image composition (or splicing) is one of the most common image manipulation operations. One such example is shown in Fig. 1, in which the girl on the right is inserted. Although this image shows a harmless manipulation case, several more controversial cases have been reported.

In this spirit, Riess and Angelopoulou [2] proposed to analyze illuminant color estimates from local image regions. Unfortunately, the interpretation of their resulting so-called illuminant maps is left to

human experts. As it turns out, this decision is, in practice, often challenging. Moreover, relying on visual assessment can be misleading, as the human visual system is quite inept at judging illumination environments in pictures [3]. Thus, it is preferable to transfer the tampering decision to an objective algorithm. In this work, we make an important step towards minimizing user interaction for an illuminant-based
tampering decision-making. We propose a new automatic method that is also significantly more reliable than earlier approaches. Quantitative evaluation shows that the proposed method achieves a detection rate of 86%, while existing illumination-based work is slightly better than guessing. We exploit the fact that local illuminant estimates are most discriminative when comparing objects of the same (or similar) material. Thus, we focus on the automated comparison of human skin, and more specifically faces, to classify the illumination on a pair of faces as either consistent or inconsistent.
In summary, the main contributions of this work are:

- Interpretation of the illumination distribution as object texture for feature computation.
- A novel edge-based characterization method for illuminant maps which explores edge attributes related to the illumination process.
- The creation of a benchmark dataset comprised of 100 skilfully created forgeries and 100 original photographs.

In Section II, we briefly review related work in color constancy and illumination-based detection of image splicing. In Section III, we present examples of illuminant maps and highlight the challenges in their exploitation. An overview of the proposed methodology, followed by a detailed explanation of all the algorithmic steps is given in Section IV. In Section V, we introduce the proposed benchmark database and present experimental results. Conclusions and potential future work are outlined in Section VI.

II. Related work

Illumination-based methods for forgery detection are either geometry-based or color-based. Geometry-based methods focus at detecting inconsistencies in light source positions between specific objects in the scene [5]–[11]. Color-based methods search for inconsistencies in the interactions between object color and light color [2], [12], [13] direction. To achieve this, a 3- Two methods have been proposed that use the direction of the incident light for exposing digital forgeries. Riess and Angelopoulou [2] followed a different approach by using a physics-based color constancy algorithm that operates on partially specular pixels. In this approach, the automatic detection of highly specular regions is avoided. The authors propose to

segment the image to estimate the illuminant color locally per segment. Recoloring each image region according to its local illuminant estimate yields a so-called illuminant map .Implausible illuminant color estimates point towards a manipulated region. Unfortunately, the authors do not provide a numerical decision criterion for tampering detection. Thus, an expert is left with the difficult task of visually examining an illuminant map for evidence of tampering. The involved challenges are further discussed in Section
In the field of color constancy, descriptors for the illuminant color have been extensively studied. Most research in color constancy focuses on uniformly illuminated

scenes containing a single dominant illuminant. However, in order to use the color of the incident illumination as a sign of image tampering, we require multiple, spatially-bound illuminant estimates. So far, limited research has been done in this direction. The work by Bleier et al. indicates that many off-the-shelf single-illuminant algorithms do not scale well on smaller image regions. Thus, problem-specific illuminant estimators are required.

Gijsenij et al. [21] proposed a pixel wise illuminant estimator. It allows to segment an image into regions illuminated by distinct illuminates. Differently illuminated regions can have crisp transitions, for instance between sunlit and shadow areas. While this is an interesting approach. a single illuminant estimator can always fail.

Fig. 2. Example illuminant map that directly shows an inconsistency.

Thus, for forensic purposes, we prefer a scheme that combines the results of multiple illuminant estimators. Earlier, Kawakami et al. [22] proposed a physics-based approach that is custom-tailored for discriminating shadow/sunlit regions. However, for our work, we consider the restriction to outdoor images overly limiting. In this paper, we build upon the ideas by [2] and [13]. We use the relatively rich illumination information provided by both physics-based and statistics-based color constancy methods as in [2], [23]. Decisions with respect to the illuminant color estimators are completely taken away from the user, which differentiates this paper from prior work.

III. Challenges in exploiting illuminant maps

To illustrate the challenges of directly exploiting illuminant estimates, we briefly examine the illuminant maps generated by the method of Riess and Angelopoulou [2]. In this approach, an image is subdivided into regions of similar color (super pixels). An illuminant color is locally estimated using the pixels within each superpixel (for details, see [2] and Section IV-A). Recoloring each superpixel with its local illuminant color estimate yields a so- called illuminant map. A human expert can then investigate the input image and the illuminant map to detect inconsistencies.

Fig. 2 shows an example image and its illuminant map, in which an inconsistency can be directly shown: the inserted mandarin orange in the top right exhibits multiple green spots in the illuminant map. All other fruits in the scene show a gradual transition from red to blue. The inserted mandarin orange is the only one that deviates from this pattern.

In practice, however, such analysis is often challenging, as shown in Fig. 3. The top left image is original, while the bottom image is a composite with the right-most girl inserted. Several illuminant estimates are clear outliers, such as the hair of the girl on the left in the bottom image, which is estimated as strongly red illuminated. Thus, from an expert's viewpoint, it is reasonable to discard such regions and to focus on more reliable regions, e.g., the faces. In Fig. 3, however, it is difficult to justify a tampering decision by comparing the color distributions in the facial regions. It is also challenging to argue, based on these illuminant maps, that the right-most girl in the bottom image has been inserted, while, e.g., the right-most boy in the top image is

original. Christo Ananth et al. [4] proposed a system which uses intermediate features of maximum overlap wavelet transform (IMOWT) as a pre-processing step. The coefficients derived from IMOWT are subjected to 2D histogram Grouping. This method is simple, fast and unsupervised. 2D histograms are used to obtain Grouping of color image. This Grouping output gives three segmentation maps which are fused together to get the final segmented output. This method produces good segmentation results when compared to the direct application of 2D Histogram Grouping. IMOWT is the efficient transform in which a set of wavelet features of the same size of various levels of resolutions and different local window sizes for different levels are used. IMOWT is efficient because of its time effectiveness, flexibility and translation invariance which are useful for good segmentation results.

Fig. 3. Example illuminant maps for an original image (top) and a spliced image (bottom). The illuminant maps are created with the IIC-based illuminant estimator (see Section IV-A).

The features are designed to capture the shape of the superpixels in conjunction with the color distribution. In this spirit, our goal is to replace the expert-in-the-loop, y only requiring annotations of faces in the image. Note that, the estimation of the illuminant color is error-prone and affected by the materials in the scene. However, (cf. Also Fig. 2), estimates on objects of similar material exhibit a lower relative error. Thus, we limit our detector to skin, and in particular to faces. Pigmentation is the most obvious difference in skin characteristics

IV. Overview and algorithmic details

We classify the illumination for each pair of faces in the image as either consistent or inconsistent. Throughout the paper, we abbreviate illuminant estimation as IE, and illuminant maps as IM. The proposed method consists of five main components:

1) Dense Local Illuminant Estimation (IE): The input image is segmented into homogeneous regions. Per illuminant estimator, a new image is created where each region is colored with the extracted illuminant color. This resulting intermediate representation is called illuminant map (IM).

2) Face Extraction: An automated face detector can be employed. We then crop every bounding box out of each illuminant map, so that only the illuminant estimates of the face regions remain.

3) Computation of Illuminant Features: for all face regions, texture-based and gradient-based features are computed on the IM values. Each one of them encodes complementary information for classification.

4) Paired Face Features: Our goal is to assess whether a pair of faces in an image is consistently illuminated. For an image with n_f faces, we construct joint feature vectors, consisting of all possible pairs of faces.

5) Classification: We use a machine learning approach to automatically classify the feature vectors. We consider an image as a forgery if at least one pair of faces in the

image is classified as inconsistently illuminated.

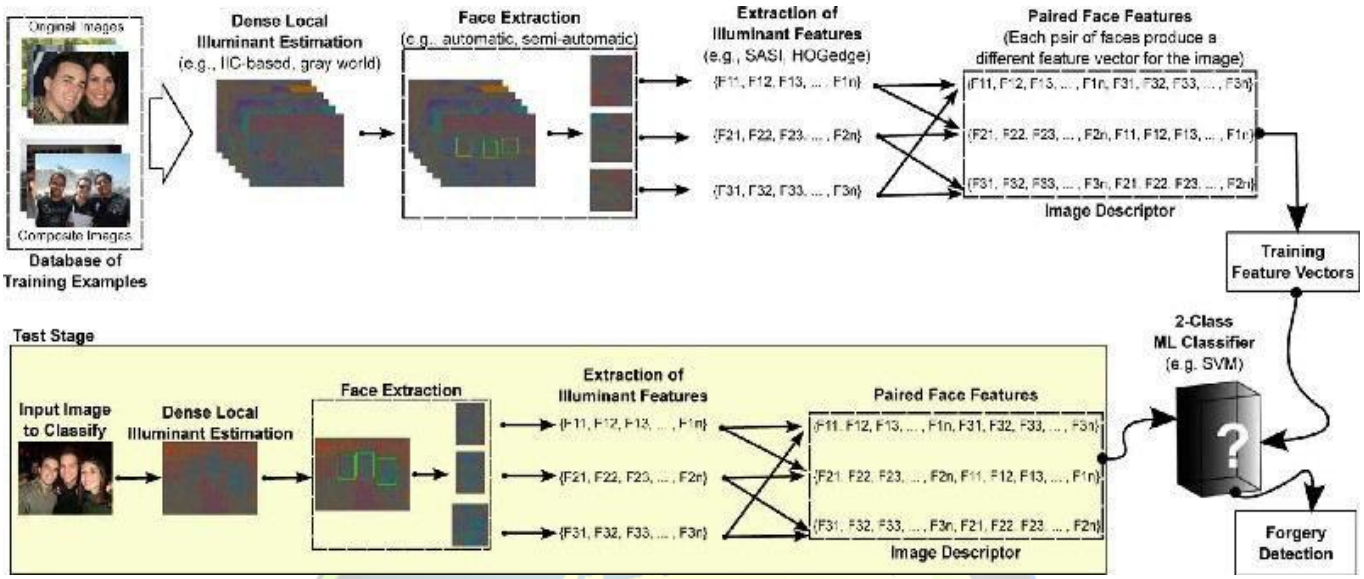

Fig. 4. Overview of the proposed method.

Fig. 4 summarizes these steps. In the remainder of this section, we present the details of these components.

A. Dense Local Illuminant Estimation

To compute a dense set of localized illuminant color estimates, the input image is segmented into super pixels, i.e., regions of approximately constant chromaticity, using the algorithm by Felzenszwalb and Huttenlocher [25]. Per superpixel, the color of the illuminant is estimated. We use illuminant color estimators: the statistical generalized gray world estimates. We obtain illuminant maps by recoloring each superpixel with the estimated illuminant chromaticities of each one of the estimators.

Generalized Gray World Estimates: The classical gray world assumption by Buchsbaum [26] states that the average color of a scene is gray.

Thus, a deviation of the average of the image intensities from the expected gray color is due to the illuminant. Although this assumption is nowadays considered to be overly simplified [17], it has inspired the further design of statistical descriptors for color constancy. We follow an extension of this idea, the generalized gray world approach by van de Weijer et al. [23].

Let $f(X)=(f_R(X),f_G(X),f_B(X))^T$ denote the observed RGB color of a pixel at location . Van deWeijer et al. s[23] assume purely diffuse reflection and linear camera response. Then, f(X) is formed by

$$f(X)=\int_\Omega e(\lambda,x)s(\lambda,x)c(\lambda)d\lambda,$$

Where Ω denotes the spectrum of visible light, λ denotes the wavelength of the light, e(λ,x) denotes the spectrum of the illuminant, s(λ,x) the surface reflectance of an object, and c(λ) the color sensitivities of the camera (i.e., one function per color channel).

B. Face Extraction

We require bounding boxes around all faces in an image that should be part of the investigation. For obtaining the bounding boxes, we could in principle use an automated algorithm, e.g., the one by Schwartz et al. [30].

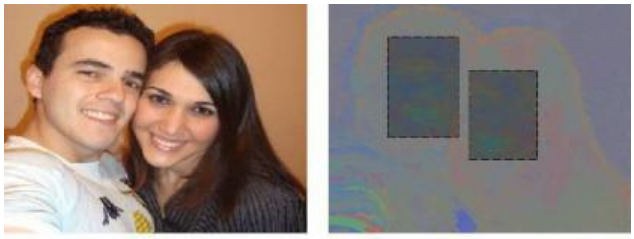

(a) (b)

Fig. 5. Original image and its gray world map. Highlighted regions in the gray world map show a similar appearance. (a) Original. (b) Gray world with highlighted similar parts.

For instance, consider an image where all persons of interest are illuminated by flashlight. The illuminates are expected to agree with one another. Conversely, assume that a person in the foreground is illuminated by flashlight, and a person in the background is illuminated by ambient light. Then, a difference in the color of the illuminates is expected. Such differences are hard to distinguish in a fully-automated manner, but can be easily excluded in manual annotation.

We illustrate this setup in Fig. 5. The faces in Fig. 5(a) can be assumed to be exposed to the same illuminant. As Fig. 5(b) shows, the corresponding gray world illuminant map for these two faces also has similar values. The parameter values were previously investigated by Riess and Angelopoulou [2], [29]. In this paper, we rely on their findings.

C. Texture Description: SASI Algorithm

We use the Statistical Analysis of Structural Information (SASI) descriptor by Carkacioglu and Yarman-Vural [31] to extract texture information from illuminant maps. Recently, Penatti et al. [32] pointed out that SASI performs remarkably well. For our application, the most important advantage of SASI is its capability of capturing small granularities and discontinuities in texture patterns. Distinct illuminant colors interact differently with the underlying surfaces, thus generating distinct illumination "texture". This can be a very fine texture, whose subtleties are best captured by SASI.

SASI is a generic descriptor that measures the structural properties of textures. It is based on the autocorrelation of horizontal, vertical and diagonal pixel lines over an image at different scales. Instead of computing the autocorrelation for every possible shift, only a small number of shifts is considered. One autocorrelation is computed using a specific fixed orientation, scale, and shift.

Computing the mean and standard deviation of all such pixel values yields two feature dimensions.

Repeating this computation for varying orientations, scales and shifts yields a 128-dimensional feature vector. As a final step, this vector is normalized by subtracting its mean value, and dividing it by its standard deviation. For details, please refer to [31].

D. Interpretation of Illuminant Edges: Hogedge Algorithm

Differing illuminant estimates in neighbouring segments can lead to discontinuities in the illuminant map. Dissimilar illuminant estimates can occur for a number of reasons: changing geometry, changing material, noise, retouching or changes in the incident light. Thus, one can interpret an illuminant estimate as a low-level descriptor of the underlying image statistics. We observed that the edges, e.g., computed by a Canny edge detector, detect in several cases a combination of the segment borders and isophotes (i.e., areas of similar incident light in the image). When an image is spliced, the statistics of these edges is likely to differ from original images. To characterize such edge discontinuities, we propose a new feature descriptor called HOGedge. It is based on the well-known HOG-descriptor, and computes visual dictionaries of gradient intensities in edge points. The full algorithm is described in the remainder of

this section. Fig. 6shows an algorithmic overview of the method. We first extract approximately equally distributed candidate points on the edges of illuminant maps. At these points, HOG descriptors are computed.

These descriptors are summarized in a visual words dictionary. Each of these steps is presented in greater detail in the next subsections.

Extraction of Edge Points: Given a face region from an illuminant map, we first extract edge points using the Canny edge detector [33]. This yields a large number of spatially close edge points. To reduce the number of points, we filter the Canny output using the following rule: starting from a seed point, we eliminate all other edge pixels in a region of interest (ROI) centred around the seed point. The edge points that are closest to the ROI (but outside of it) are chosen as seed points for the next iteration. By iterating this process over the entire image, we reduce the number of points but still ensure that every face has a comparable density of points.

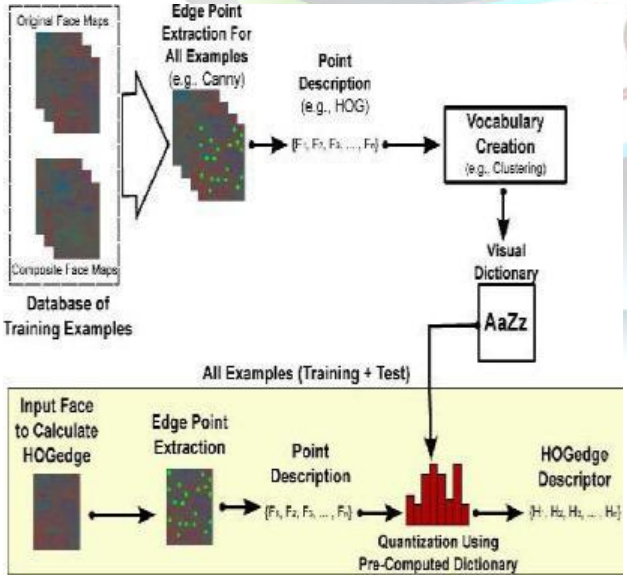

Fig. 6. Overview of the proposed H OGedge algorithm.

Fig. 7 depicts an example of the resulting points.

Point Description: We compute Histograms of Oriented Gradients (HOG) [34] to describe the distribution of the selected edge points. HOG is based on normalized local histograms of image gradient orientations in a dense grid. The HOG descriptor is constructed around each of the edge points. The neighbourhood of such an edge point is called a cell. Each cell provides a local 1-D histogram of quantized gradient directions using all cell pixels. To construct the feature vector, the histograms of all cells within a spatially larger region are combined and contrast-

normalized. We use the HOG output as a feature vector for the subsequent steps.

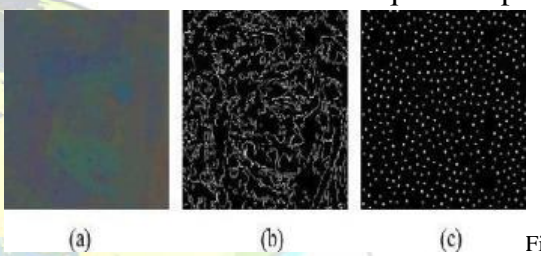

Fig. 7.

(a) Gray world IM for the left face in Fig. 6(a). (b) Result of the Canny edge detector when applied on this IM. (c) Final edge points after filtering using a square region. (a) IM derived from gray world. (b) Canny Edges. (c) Filtered Points.

Visual Vocabulary: The number of extracted HOG vectors varies depending on the size and structure of the face under examination. We use visual dictionaries [35] to obtain feature vectors of fixed length. Visual dictionaries constitute a robust representation, where each face is treated as a set of region descriptors. The spatial location of each region is discarded [36]. To construct our visual dictionary, we subdivide the training data into feature vectors from original and doctored images. Each group is clustered in clusters using the -means algorithm [37]. Then, a visual dictionary with visual words is constructed, where each word is represented by a cluster centre. Thus, the visual dictionary summarizes the most representative feature vectors of the

training set.

The horizontal axis corresponds to different feature dimensions, while the vertical axis represents the average feature value for different combinations of descriptors and illuminant maps. From top to bottom, left to right: SASI-IIC, HOGedge-IIC, SASI-Gray-World, HOGedge-Gray-World.

E. Face Pair

To compare two faces, we combine the same descriptors for each of the two faces. For instance, we can concatenate the SASI-descriptors that were computed on gray world.

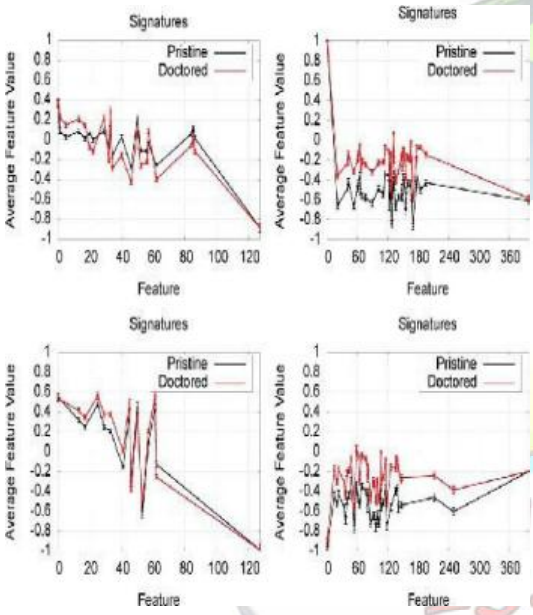

Fig. 8. Average signatures from original and spliced images.

The idea is that a feature concatenation from two faces is different when one of the faces is an original and one is spliced. For an image containing n_f faces, the number of face pairs is $(n_f(n_f-1))/2$

The SASI and HOGedge descriptors capture two different properties of the face regions. From a signal processing point of view, both descriptors are signatures with different behaviour. Fig. 8 shows a very high-level visualization of the distinct information that is captured by these two descriptors. For one of the folds of our experiments (see Section V-C), we computed the mean value and standard deviation per feature dimension. For a less cluttered plot, we only visualize the feature dimensions with the largest difference in the mean values for this fold. This experiment empirically demonstrates two points. Firstly, SASI and HOGedge, in combination with the IIC-based and gray world illuminant maps create features that discriminate well between original and tampered images, in at least some dimensions. Secondly, the dimensions, where these features have distinct value, vary between the four combinations of the feature vectors. We exploit this property during classification by fusing the output of the classification on both feature sets, as described in the next section.

F. Classification

We classify the illumination for each pair of faces in an image as either consistent or inconsistent. Assuming all selected faces are illuminated by the same light source, we tag an image as manipulated if one pair is classified as inconsistent. Individual feature vectors, i.e., SASI or HOGedge features on either gray world or IIC-based illuminant maps, are classified using a support vector machine (SVM) classifier with a radial basis function (RBF) kernel.

The information provided by the SASI features is complementary to the information from the HOGedge features. Thus, we use a machine learning-based fusion technique for improving the detection performance. Inspired by the work of Ludwig et al. [38], we use a late fusion technique named SVM-Meta Fusion. We classify each combination of illuminant map and feature type independently (i.e., SASI-Gray-World, SASI-IIC,HOGedge-Gray-World and HOGedge-IIC) using a two-class SVM classifier to obtain the distance between the image's feature vectors and the classifier decision boundary. SVM-Meta Fusion then

merges the marginal distances provided by all individual classifiers to build a new feature vector. Another SVM classifier.

V. Conclusions and future work

In this work, we presented a new method for detecting forged images of people using the illuminant color. We estimate the illuminant color using a statistical gray edge method and a physics-based method which exploits the inverse intensity- chromaticity color space. We treat these illuminant maps as texture maps. We also extract information on the distribution of edges on these maps. In order to describe the edge information, we propose a new algorithm based on edge-points and the HOG descriptor, called HOGedge. We combine these complementary cues (texture- and edge-based) using machine learning late fusion. Our results are encouraging, yielding an AUC of over 86% correct classification. Good results are also achieved over internet images and under cross-database training/testing. Although the proposed method is custom-tailored to detect splicing on images containing faces, there is no principal hindrance in applying it to other, problem-specific materials in the scene.

The proposed method requires only a minimum amount of human interaction and provides a crisp statement on the authenticity of the image. Additionally, it is a significant advancement in the exploitation of illuminant color as a forensic cue. Prior color-based work either assumes complex user interaction or imposes very limiting assumptions. Although promising as forensic evidence; methods that operate on illuminant color are inherently prone to estimation errors. Thus, we expect that further improvements can be achieved when more advanced illuminant color estimators become available. For instance, while we were developing this work, Bianco and Schettini [49] proposed a machine-learning based illuminant estimator particularly for faces. An incorporation of this method is subject of future work.

Reasonably effective skin detection methods have been presented in the computer vision literature in the past years. Incorporating such techniques can further expand the applicability of our method. Such an improvement could be employed, for instance, in detecting pornography compositions which, according to forensic practitioners, have become increasingly common nowadays.

REFERENCES

[1] A. Rocha, W. Scheirer, T. E. Boult, and S. Goldenstein, "Vision of the unseen: Current trends and challenges in digital image and video forensics," ACM Comput. Surveys, vol. 43, pp. 1–42, 2011.

[2] C. Riess and E. Angelopoulou, "Scene illumination as an indicator of image manipulation," Inf. Hiding, vol. 6387, pp. 66–80, 2010.

[3] H. Farid and M. J. Bravo, "Image forensic analyses that elude the human visual system," in Proc. Symp. Electron. Imaging (SPIE), 2010, pp. 1–1

[4] Christo Ananth, A.S.Senthilkani, S.Kamala Gomathy, J.Arockia Renilda, G.Blesslin Jebitha, Sankari @Saranya.S., "Color Image Segmentation using IMOWT with 2D Histogram Grouping", International Journal of Computer Science and Mobile Computing (IJCSMC), Vol. 3, Issue. 5, May 2014, pp-1 – 7

[5] H. Farid, A 3-D lighting and shadow analysis of the JFK Zapruder film (Frame 317), Dartmouth College, Tech. Rep. TR2010–677, 2010.

[6] M. Johnson and H. Farid, "Exposing digital forgeries by detecting inconsistencies in lighting," in Proc. ACM Workshop on Multimedia and Security, New York, NY, USA, 2005, pp. 1–10.

[7] M. Johnson and H. Farid, "Exposing digital forgeries in complex lighting environments," IEEE Trans. Inf. Forensics Security, vol. 3, no. 2, pp. 450–461, Jun. 2007.

[8] M. Johnson and H. Farid, "Exposing digital

forgeries through specular highlights on the eye," in Proc. Int.Workshop on Inform. Hiding, 2007, pp. 311–325.
[9] E. Kee and H. Farid, "Exposing digital forgeries from 3-D lighting environments," in Proc. IEEE Int. Workshop on Inform. Forensics and Security (WIFS), Dec. 2010, pp. 1–6.
[10] W. Fan, K. Wang, F. Cayre, and Z. Xiong, "3D lighting-based image forgery detection using shape-from-shading," in Proc. Eur. Signal Processing Conf. (EUSIPCO), Aug. 2012, pp. 1777–1781.
[11] J. F. O'Brien and H. Farid, "Exposing photo manipulation with inconsistent reflections," ACM Trans. Graphics, vol. 31, no. 1, pp. 1–11, Jan. 2012.
[12] S. Gholap and P. K. Bora, "Illuminant colour based image forensics," in Proc. IEEE Region 10 Conf., 2008, pp. 1–5.
[13] X.Wu and Z. Fang, "Image splicing detection using illuminant color inconsistency," in Proc. IEEE Int. Conf. Multimedia Inform. Networking and Security, Nov. 2011, pp. 600–603.
[14] P. Saboia, T. Carvalho, and A. Rocha, "Eye specular highlights telltales for digital forensics: A machine learning approach," in Proc. IEEE Int. Conf. Image Processing (ICIP), 2011, pp. 1937–1940.
[15] C. Riess and E. Angelopoulou, "Physics-based illuminant color estimation as an image semantics clue," in Proc. IEEE Int. Conf. Image Processing, Nov. 2009, pp. 689–692.
[16] K. Barnard, V. Cardei, and B. Funt, "A comparison of computational color constancy algorithms–Part I: Methodology and Experiments With Synthesized Data," IEEE Trans. Image Process., vol. 11, no. 9, pp. 972–983, Sep. 2002.
[17] K. Barnard, L. Martin, A. Coath, and B. Funt, "A comparison of computational color constancy algorithms – Part II: Experiments With Image Data," IEEE Trans. Image Process., vol. 11, no. 9, pp. 985–996, Sep. 2002.
[18] A. Gijsenij, T. Gevers, and J. van deWeijer, "Computational color constancy: Survey and experiments," IEEE Trans. Image Process., vol. 20, no. 9, pp. 2475–2489, Sep. 2011. [19] M. Bleier, C. Riess, S. Beigpour, E. Eibenberger, E. Angelopoulou, T. Tröger, and A. Kaup, "Color constancy and non-uniform illumination: Can existing algorithms work?," in Proc. IEEE Color and Photometry in Comput. Vision Workshop, 2011, pp. 774–781.
[20] M. Ebner, "Color constancy using local color shifts," in Proc. Eur. Conf. Comput. Vision, 2004, pp. 276–287.
[21] A. Gijsenij, R. Lu, and T. Gevers, "Color constancy for multiple light sources," IEEE Trans. Image Process., vol. 21, no. 2, pp. 697–707, Feb. 2012.
[22] R. Kawakami, K. Ikeuchi, and R. T. Tan, "Consistent surface color for texturing large objects in outdoor scenes," in Proc. IEEE Int. Conf. Comput. Vision, 2005, pp. 1200–1207.
[23] J. van de Weijer, T. Gevers, and A. Gijsenij, "Edge-based color constancy," IEEE Trans. Image Process., vol. 16, no. 9, pp. 2207–2214, Sep. 2007.
[24] T. Igarashi, K. Nishino, and S. K. Nayar, "The appearance of human skin: A survey," Found. Trends Comput. Graph. Vis., vol. 3, no. 1, pp. 1–95, 2007.
[25] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," Int. J. Comput. Vis., vol. 59, no. 2, pp. 167–181, 2004.
[26] G. Buchsbaum, "A spatial processor model for color perception," J. Franklin Inst., vol. 310, no. 1, pp. 1–26, Jul. 1980.