

Camera Response Functions for Image Forensics: An Automatic Algorithm for Splicing Detection

Yu-Feng Hsu and Shih-Fu Chang, *Fellow, IEEE*

Abstract—We present a fully automatic method to detect doctored digital images. Our method is based on a rigorous consistency checking principle of physical characteristics among different arbitrarily shaped image regions. In this paper, we specifically study the camera response function (CRF), a fundamental property in cameras mapping input irradiance to output image intensity. A test image is first automatically segmented into distinct arbitrarily shaped regions. One CRF is estimated from each region using geometric invariants from locally planar irradiance points (LPIPs). To classify a boundary segment between two regions as authentic or spliced, CRF-based cross fitting and local image features are computed and fed to statistical classifiers. Such segment level scores are further fused to infer the image level authenticity. Tests on two data sets reach performance levels of 70% precision and 70% recall, showing promising potential for real-world applications. Moreover, we examine individual features and discover the key factor in splicing detection. Our experiments show that the anomaly introduced around splicing boundaries plays the major role in detecting splicing. Such finding is important for designing effective and efficient solutions to image splicing detection.

Index Terms—Camera response function (CRF), image forensics, splicing detection, tampering detection.

I. INTRODUCTION

WITH the ease of digital image manipulation, like copy-and-paste (splicing), image forgery has become a common concern and verification of content integrity has become increasingly important. An intuitive and promising approach to image forgery detection is to examine the consistency of inherent physics-based attributes among different parts of an image. These attributes can be natural-scene related, for example, lighting, shadow, and object geometry; or they can be imaging device properties such as camera response function (CRF), demosaicking filter, and sensor noise statistics. Any image that fails to show consistency in these two aspects may be considered as suspects of forgery. Such approaches are

passive—no active mechanisms are needed to generate and embed additional signatures into images.

Recent research in computer vision and image processing has studied various physical characteristics of imaging devices. A typical camera imaging pipeline is illustrated in Fig. 1. The light rays (*radiance*s) are first refracted through the optical lens and recorded onto the CCD (or CMOS) sensor. Since the CCD sensor only records a single color channel signal, a demosaicking color filter is needed to interpolate the CCD output into a multispectral image (the *irradiance*, usually in RGB or CMYK color spaces). The following component, CRF, transforms the interpolated irradiance signal nonlinearly to compensate for the measurable dynamic range, producing the final output, denoted as *brightness* or *intensity*. Components in this pipeline, such as the CCD sensor, demosaicking filter, and the CRF, may possess unique characteristics to camera models or even to camera units. If successfully extracted, these physical attributes can be used as the “fingerprint” of the capturing source. Some prior works have proposed fingerprint recovery methods for image source identification, including demosaicking filter [1]–[3], CCD sensor noise [4], CRF [5]–[9] and inherent relations between different color channels within the camera [10]. They have also led to consistency checking schemes based on demosaicking inconsistency [1]–[3], CCD sensor noise [11], CRF abnormality [12], and CRF inconsistency [13].

Common inconsistency checking usually estimates the most likely source model given a test image, and checks whether there is inconsistency between any parts and the overall estimated model [1], [3], [11]. In this paper, on the other hand, we address the more general question whether physical attributes from two image parts are from the same source without needing to estimate the overall model. This may be formulated as hypothesis testing or discriminative classification, requiring appropriate similarity measures [13]–[15].

We focus on camera-signature-based inconsistency measures in this paper, particularly CRF. Among all available signatures, CRF has been extensively studied with many robust estimation methods proposed [5]–[8], [16]–[19]. However, the space of CRFs has been observed to be extremely dense: CRFs from different cameras can appear very close to each other. Therefore, the consistency measure computed directly from the CRFs would not be effective for splicing detection. One needs more relevant physical cues, e.g., splicing-induced anomalies.

In this paper, we introduce an automatic image splicing detection method based on CRF consistency checking. The objective is to construct a system free of human labor without severely compromising its detection power. An overview of our approach is given in Fig. 2. We run CRF estimation on each of the automatically segmented image regions, construct a feature

Manuscript received December 29, 2009; revised September 04, 2010; accepted September 08, 2010. Date of publication September 20, 2010; date of current version November 17, 2010. This work was supported by the NSF Cyber Trust program under awards IIS-04-30258 and CNS-07-16203. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Min Wu.

Y.-F. Hsu is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA and also with Information and Communications Research Laboratories, Industrial Technology Research Institute, Chungtung Township, Hsinchu County, 310, Taiwan (e-mail: yfhsu@ee.columbia.edu).

S.-F. Chang is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: sfchang@ee.columbia.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIFS.2010.2077628

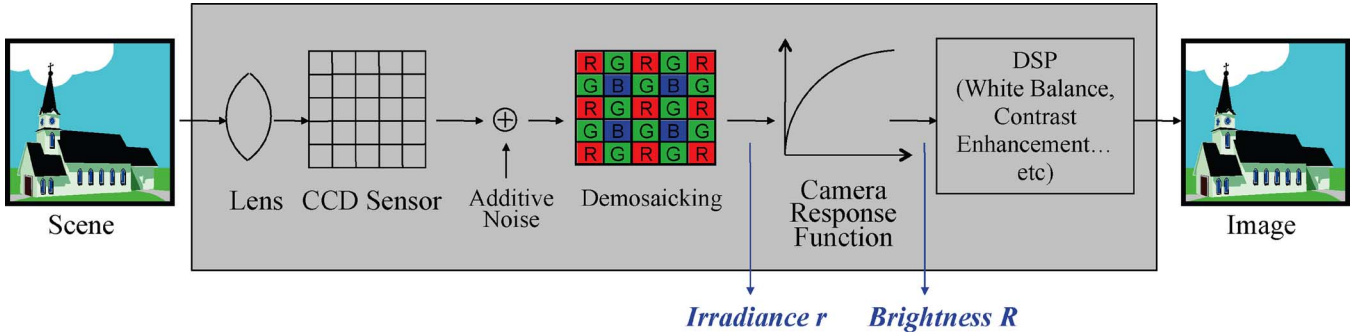


Fig. 1. Camera imaging pipeline.




Fig. 2. Automatic system for local spliced region detection.

vector for each boundary segment, and train a statistical classifier to classify these segments as authentic or spliced. We then fuse segment level results to obtain image level decisions.

We evaluate the performance on both benchmark and challenging data sets. We further perform a series of feature selection experiments to discover the main contributing factor in our detection system. The importance of the anomaly introduced near the splicing boundary has been consistently revealed and has led us to a new, effective feature set which indeed improves the accuracy significantly.

The contributions of this paper are multifold. First, we develop a fully automatic system with an aim to locate any arbitrarily shaped spliced segment, not restricted to fixed, coarse blocks. Second, through extensive feature selection, we discover the splicing boundary anomaly more crucial than pair-wise attribute difference. Lastly, we develop a theoretical model to predict and support the fusion results from local segment scores to image level decisions.

Our method conceptually is most related to the work in [12], in which suspicious boundaries are manually identified and anomalies associated with the CRF estimated from these selected segments are used to detect splicing. There are interesting similarities and differences. Our results in feature selection confirm that abnormal boundary segments created by splicing indeed provide the most beneficial features for tampering detection, justifying the focus on the boundaries both in our work and [12]. However, our method is different in our rich information from CRF cross fitting using locally planar irradiance points (LPIPs) from segmented regions instead of the anomalies only from the estimated CRF, as in [12]. Our algorithm is also fully automatic, while [12] requires manual selection of suspicious boundaries. Finally, we work from segment level anomaly detection and fuse the scores to image level classification results, a different methodology than [12].

When compared to the double quantization detection in [14], we are able to detect arbitrarily shaped regions instead of fixed

size blocks. Also, while they determine the authenticity of individual blocks, we focus on the relations between two distinct areas. In fact, these two types of features are complementary. In our work reported in [20] and [21], we have explored the integration of these two distinct features and demonstrated higher tampering detection accuracy.

This paper is organized as follows: Section II introduces the CRF with its mathematical models and estimation methods. Section III describes our automatic splicing detection system followed by results in Section IV. Section V includes theoretical analysis of the fusion model. The feature selection study is then presented in Section VI with results in Section VII. Finally, Section VIII concludes this paper.

II. CAMERA RESPONSE FUNCTION

Prior to the image output recorded on memory cards, cameras perform various operations on the incoming light rays: linear or nonlinear, point-wise or spatial, all of which when combined yield visually pleasant images to human eyes. CRF, the most salient point-wise operation, maps scene irradiance to image brightness nonlinearly (Fig. 1) [19]. Though some emerging camera models may add spatially varying CRFs, it is still considered valid to assume one invariant CRF per camera for existing models today. Mathematically, CRF is often denoted as a single-variable function $R = f(r)$. Without loss of generality, both r and R are assumed to be between $[0, 1]$ although different manufacturers may produce different dynamic ranges.

Several models have been proposed to represent CRF: one-parameter gamma function, $R = f(r) = r^{\alpha_0}$ [6], polynomial $R = f(r) = \sum_{n=0}^N r^{\beta_n}$ [18], PCA-based empirical model of response (EMOR) [19], and the generalized gamma curve model (GGCM) $R = f(r) = r \sum_{i=0}^n \alpha_i r^i$ [8]. A comparison across these models is given in [8], showing superior approximation of EMOR and GGCM over polynomial models. Considering the tradeoff between modeling power and computational complexity, we use a first-order GGCM, $n = 1$.

A. CRF Estimation

The standard CRF calibration procedure is carried out using Macbeth Charts with known irradiances [7]–[9]. The CRF is obtained by fitting a curve to the scatterplot of measured brightness values against their corresponding irradiances. Such CRFs would serve as ground truth when automatic estimation algorithms are evaluated.

The interest of CRF estimation from real images has spanned over a decade. Earlier efforts used images of the same scene with multiple known exposures and solved for the CRF under over-constrained conditions [16], [17]. Single-image CRF estimation algorithms are only proposed recently. The method in [6] looks for the optimal gamma function with the least higher order correlation. The method in [7] uses RGB colinearity resulted from color blending on edges, which is further extended to grayscale images in [9]. The algorithm in [8] is based on the observation that the ratios between partial derivatives at LPIPs carry information solely related to the CRF but not to the image content. Therefore, by applying an initial step to extract qualifying LPIPs, the CRF can be successfully recovered. This algorithm is discussed with more details in Section II-B and will be used in our splicing detection system.

B. CRF Estimation Using Geometry Invariants

In [8], the partial derivatives in the brightness domain are shown to be related to the CRF and the irradiance information

$$\begin{aligned} R_x &= f'(r)r_x, R_y = f'(r)r_y \\ R_{xx} &= f''(r)r_x^2 + f'(r)r_{xx} \\ R_{yy} &= f''(r)r_y^2 + f'(r)r_{yy} \\ R_{xy} &= f''(r)r_xr_y + f'(r)r_{xy} \end{aligned}$$

where $r_x, r_y, r_{xx}, r_{yy}, r_{xy}$ denote partial spatial derivatives of the unknown geometry of irradiance. If a point has a locally planar irradiance geometry, $r = ax + by + c$, then the second-order partial derivatives in the irradiance domain r_{xx}, r_{yy}, r_{xy} would all be zero, and the following equation holds:

$$\frac{R_{xx}}{R_x^2} = \frac{R_{xy}}{R_xr_y} = \frac{R_{yy}}{R_y^2} = \frac{f''(r)}{(f'(r))^2} = \frac{f''(f^{-1}(R))}{(f'(f^{-1}(R)))^2} \quad (1)$$

This quantity, denoted as $A(R)$, does not carry any information about the geometry of r , i.e., $\{a, b, c\}$. With further manipulation we get another quantity $Q(R)$

$$\text{Geometry Invariant : } Q(R) = \frac{1}{1 - A(R)R} \quad (2)$$

which is also independent of irradiance geometry and hence termed Geometry Invariant (GI). It is equal to the gamma parameter α_0 if the CRF takes the gamma form. When the CRF is represented by a first-order GGCM, $Q(R)$ is related to the parameters $\{\alpha_0, \alpha_1\}$ by

$$Q(R) = \frac{(\alpha_1 r \ln(r) + \alpha_1 r + \alpha_0)^2}{\alpha_0 - \alpha_1 r}. \quad (3)$$

Given a single image, the CRF is estimated by first extracting LPIPs, computing GIs (i.e., Q and R values) using (1) and (2), and iteratively looking for the optimal GGCM parameters to fit the computed GI values using (3). The authors of [8] further

explored rigorous ways of LPIP selection, error metric definition, and cross color channel similarity to enhance the accuracy of CRF estimation. Extensive experiments over multiple cameras and diverse images have shown excellent estimation accuracy—average root mean square error (RMSE) as low as 0.0224. More details can be found in [8].

III. AUTOMATIC SPLICING DETECTION BY CRF CONSISTENCY CHECKING

In this section, we discuss the components in our automatic splicing detection system.

Consistency checking is motivated by the fact that spliced images typically contain suspicious regions with different device signatures from other regions. If the spliced region can be automatically located and the signature correctly recovered, image tampering can be revealed successfully. In most scenarios, manual input is not available and an automated process is demanded. This brings image segmentation into the picture, followed by crucial tasks of device signature extraction and consistency checking. In this work, we choose a popular segmentation tool, Normalized Cuts [22], though other methods such as Mean Shift [23] may also be considered. For device signature extraction, we use the single-channel CRF estimation described in Section II-B. The model-based fitting in our prior work [13], [14] is used as consistency measures.

Instead of checking between every possible pair of segmented regions, we focus on the authenticity of boundary segments. The anomalous points created near splicing boundaries are expected to be crucial. Given the segmentation results of an image, we propose to first detect suspicious segments and infer if the whole image contains any spliced content.

As illustrated in Fig. 2, the test image is first segmented into several regions with qualifying LPIPs extracted and CRFs estimated. To determine the authenticity of a boundary segment, we apply cross fitting between two neighboring regions surrounding the segment and self fitting data samples from the boundary segment. The fitting scores, along with other features extracted from the regions, are used as the feature vector representing the segment under question. The segment is detected as authentic or spliced using a discriminative classifier such as support vector machine (SVM). Additionally, image level decisions are obtained by fusing segment level classification results.

A. Image Segmentation

Among various image segmentation tools, Normalized Cuts [22] is widely used because of its intuitive formulation and robust results. It treats image pixels as vertices in a graph and their dissimilarity measures as edge weights. The output is multiple subgraphs with high similarity within each subgraph and low similarity across distinct subgraphs. Practically, Normalized Cuts requires the number of desired regions to be predetermined, typically from 2 to 20. As most spliced images contain only one foreground object of reasonably large size, in our experiments, we opt for a coarse setting to avoid over-segmentation and excessive computation cost. The number of regions is set to be eight based on empirical validation. Note our consistency verification framework is general and different segmentation methods or settings can be easily incorporated.



Fig. 3. (a) Different types of relations between the automatically segmented boundaries and the true splicing boundary. (b) Definitions of regions surrounding a boundary segment in our classification system. A and B are regions adjacent to the boundary while E is the region coinciding with the boundary.

Each boundary segment is categorized into one of the following cases depending on the types of its neighboring sides [Fig. 3(a)]. 1) Authentic: Both sides are from the same camera; thus the segment under consideration is authentic. 2) Spliced: Two neighboring regions are untampered but are from different cameras. In this case, the segment coincides with the splicing boundary [white dashed contour in Fig. 3(a)]. 3) Partially aligned: One or both regions contain content from two cameras. In other words, the automatic boundary segment only overlaps partially with the actual splicing boundary. The notations are illustrated in Fig. 3(b): we shall denote the two authentic sides as area A and area B, the dilated boundary segment as area E, and O for the union of areas A, B, and E. These notations will be used consistently throughout this paper.

For splicing detection purposes, there is no need to distinguish Spliced from Partially aligned cases since they both indicate the presence of the splicing operation. However, at segment level, the authenticity of partially aligned segments is difficult to define. As such, we only include well-defined Authentic and Spliced segments as our training data to learn robust two-class classifiers. Our hypothesis, as will be confirmed later by experiments, is that such classifiers will still provide reasonable detection power on Partially aligned cases.

B. CRF Consistency Measure via Cross Fitting

To check if a boundary segment is authentic or spliced, we compute cross fitting errors using the estimated CRFs and (Q, R) values of the selected LPIPs: For $i, j \in \{A, B\}$ (where i, j are indices to areas and take values on A or B)

$$\begin{aligned} \mathbf{s}_{ij} &= \left\{ s_{ij}^{(n)} \mid n \leq N_i \right\} \\ &= \left\{ \left(Q_i^{(n)}(R) - Q_j^{(\text{model})}(R) \right)^2 \mid n \leq N_i \right\} \end{aligned} \quad (4)$$

where n is the index of LPIPs and N_i the total number of LPIPs (or (Q, R) 's) from area i . If areas A and B are indeed from different cameras, we should obtain large cross fitting errors \mathbf{s}_{ij} 's. The above equation measures how well a CRF model $(Q_j^{(\text{model})})$ estimated in one area can be used to fit the points $(Q_i^{(n)})$ extracted from a different area. Plugging the CRF model from (3) in the place of $Q_j^{(\text{model})}$, we can set

$$\mathbf{s}_{ij} = \left\{ \left(Q_{i,n} - \frac{(\alpha_{1,j} r_n \ln(r_n) + \alpha_{1,j} r_n + \alpha_{0,j})^2}{\alpha_{0,j} - \alpha_{1,j} r_n} \right)^2 \mid n \leq N_i \right\}. \quad (5)$$

Likewise, for $k \in \{E, O\}$, the self fitting errors are given by

$$\begin{aligned} \mathbf{s}_{kk} &= \left\{ s_{kk}^{(n)} \mid n \leq N_k \right\} \\ &= \left\{ \left(Q_k^{(n)}(R) - Q_k^{(\text{model})}(R) \right)^2 \mid n \leq N_k \right\} \quad (6) \\ \mathbf{s}_{kk} &= \left\{ \left(Q_{k,n} - \frac{(\alpha_{1,k} r_n \ln(r_n) + \alpha_{1,k} r_n + \alpha_{0,k})^2}{\alpha_{0,k} - \alpha_{1,k} r_n} \right)^2 \mid n \leq N_k \right\}. \quad (7) \end{aligned}$$

Anomalous distributions of (Q, R) samples from areas E and O are expected if they are not from a single camera. Thus, their self fitting results \mathbf{s}_{EE} and \mathbf{s}_{OO} should exhibit distinct behaviors from those of authentic regions. We compute the first- and second-order moments of these fitting errors and construct the first set of features of a segment

$$\mathcal{F}_{1,\mu} = [\mu(\mathbf{s}_{AA}), \mu(\mathbf{s}_{BB}), \mu(\mathbf{s}_{AB}), \mu(\mathbf{s}_{BA}), \mu(\mathbf{s}_{EE}), \mu(\mathbf{s}_{OO})] \quad (8)$$

$$\mathcal{F}_{1,\sigma} = [\sigma(\mathbf{s}_{AA}), \sigma(\mathbf{s}_{BB}), \sigma(\mathbf{s}_{AB}), \sigma(\mathbf{s}_{BA}), \sigma(\mathbf{s}_{EE}), \sigma(\mathbf{s}_{OO})] \quad (9)$$

where $\mu(\cdot)$ is the mean and $\sigma(\cdot)$ the standard deviation.

By taking the first- and second-order moments, we look at the collective statistics of fitting errors rather than individual terms. As a side note, the Bayesian classifier in [8] provides an adjustable threshold so that one can control the rate of points falsely detected as LPIPs and explore the tradeoffs between sufficient LPIPs and low errors caused by false LPIPs.

In [8], it is also found that the range of image brightness (R) significantly influences the CRF estimation accuracy (Fig. 4). As the CRF specifies the relation between input irradiance r and output intensity R over the entire range, a larger coverage of R in the observed samples will naturally lead to a smaller estimation error. A small R range, on the other hand, may induce inaccurate CRF estimation and hence segment level false alarms. Therefore, we add the averages and the ranges of R 's from each area as our second feature set

$$\mathcal{F}_{2,\mu} = [\mu(\mathbf{R}_A), \mu(\mathbf{R}_B), \mu(\mathbf{R}_E), \mu(\mathbf{R}_O)] \quad (10)$$

$$\mathcal{F}_{2,\Delta} = [\Delta(\mathbf{R}_A), \Delta(\mathbf{R}_B), \Delta(\mathbf{R}_E), \Delta(\mathbf{R}_O)] \quad (11)$$

where \mathbf{R}_A contains all the R values in area A and $\Delta(\cdot)$ is the range (difference between maximum and minimum values) of R 's. Each segment is thus represented by the combined 20-dimensional feature vector from (8) to (11).

C. SVM Classification

We use an SVM as our classifier for its wide use and proven superiority. Practically, however, due to the complexity of image content, spliced segments are often significantly outnumbered by authentic ones (Table I). It implies a biased training process toward negative samples and thus poor classification accuracy. We adopt SVM bagging as the remedy. We divide the larger pool (in this case the authentic segments) into P subsets, each with a




Fig. 4. Higher CRF estimation accuracy is expected when the range of the intensity values (R) in the test image is higher.

TABLE I
NUMBERS OF TEST SEGMENTS AND IMAGES

Segments			Images	
Auth	Spic	Amb	Auth	Spic
828	219	249	84	89

similar number of samples as the smaller pool (the spliced segments) and train P classifiers from these evenly populated samples. At the test stage, every segment gets P classified labels and distances to the decision boundaries, $[l_p, d_p], p = 0 \dots P - 1$. The distances d_p 's are warped to $[0, 1]$ with a sigmoid function and linearly fused to obtain the final label

$$l_{\text{fuse}} = \text{sign} \left(\frac{1}{P} \sum_{p=0}^{P-1} \frac{1}{1 + \exp(-d_p/\omega)} - \tau \right) \quad (12)$$

where ω controls the “bandwidth” of each single distance. In our experiment, P is set to 5, and ω is determined empirically through cross-validation. The decision threshold τ (currently set at 0.5) can be changed to obtain different operating points in the precision-recall curve as reported in later sections.

D. Image Level Fusion

To obtain a global decision for the image, naively averaging over all scores of individual segments would not be appropriate, since an image with only one spliced segment is still considered as spliced, but the average score will be small if there are multiple authentic segments with low SVM scores.

We adopt a simple alternative—if one or more segments are detected as spliced with scores larger than a threshold, the image is considered spliced. This is equivalent to an OR or MAX fusion model. Varying thresholds will result in different operating points in the image level precision-recall curve.

Note more sophisticated image level detection models can be developed by incorporating the spatial relationship of individual boundary segments. For example, segments aligned with different parts of a large object are likely to have the same classification (spliced or authentic). Such topics are beyond the scope of this paper and may be addressed as future work.

IV. EVALUATION OF BASELINE SYSTEM

In this section, we present the evaluation of our proposed method over the Basic and Advanced sets. We also test the sensitivity of the method to image compression (like JPEG). The theoretical analysis of image level fusion and comparative feature selection will be addressed in later sections.

A. Data Sets

Our benchmark data set consists of 363 uncompressed images [13]: 183 authentic and 180 spliced. Authentic images are taken with four cameras: Canon G3, Nikon D70, Canon EOS 350D, and Kodak DCS330. Each spliced image has content from exactly two cameras with a salient object from one image copied and pasted onto another using Adobe Photoshop without post-processing. We also made best efforts to ensure image content diversity with different objects and scenes. This set will be referred to as the Basic data set. We use half of these images for training and another half for testing. The numbers of images used in the testing stage are listed in Table I.

The boundary segments are categorized into three sets: Authentic, Spliced, and Partially aligned, as defined in Section III-A. Categorization is done based on the percentage of its overlap with the closest splicing boundary. If the overlap exceeds a certain threshold (e.g., 90%), it is labeled as Spliced. If lower than another threshold (e.g., 10%), it is labeled as Authentic. All others will be labeled as Partially aligned (Ambiguous). As reported in Table I, within one image, there are usually as many Partially aligned segments as Spliced segments; therefore, Partially aligned cases must not be overlooked. As mentioned in Section III-A, these segments are included in the testing stage but not the training stage.

Both the CRF estimation and cross fitting computation require certain numbers of LPIPs to work well. One reasonable question then is whether there exist enough LPIPs on the spliced boundaries in practical situations. To verify this, we examined our test data set and found that even without special smoothing of filtering after splicing, most of the spliced images have at least two or more segments that are associated with more than 100 LPIPs for each segment.

Both segment level and image level results are evaluated over the Basic data set. We randomly partition the data set at image level into training and test sets so that segments from the same image will not be included in both training and test sets. In order to see how well our classifier generalizes, we also test our detector on an additional image set that include 21 authentic images and 38 high-quality spliced images using advanced vision and graphic tools courtesy of Microsoft Research Asia. These images are typically JPEG compressed, with matting or color adjustment as postprocessing in addition to copy-and-paste. It is a much more realistic and challenging set for splicing detection, hence denoted as the Advanced data set. Note the test is performed on the new advanced spliced images without any re-training of the classifiers.

B. Performance on the Basic Data Set

Previous works on tampering detection report detection rates [1], confusion matrices [2], or simply case-by-case results [11], [12]. However, we opt for the precision-recall (PR) curve due to highly unbalanced data (Table I). The precision is defined as the ratio of correctly detected spliced segments over all segments classified as spliced, and the recall is the ratio of correctly detected spliced segments over all spliced segments. The recall rate is the same metric as detection rate, and in one of the false alarm rate definitions, the precision is simply $1 - \{\text{False Alarm Rate}\}$; therefore, the PR curve is arguably

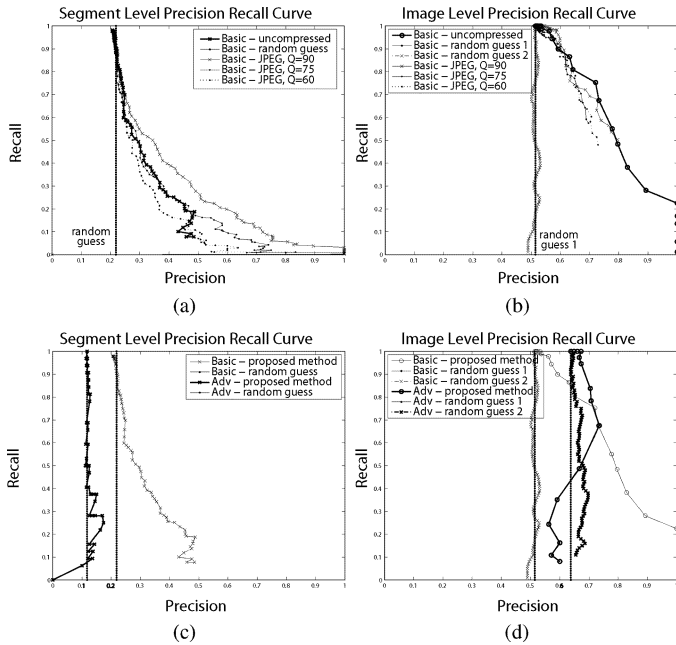


Fig. 5. PR curves for segment- and image-level splicing detection. The proposed method clearly outperforms the random-guess baseline, especially for the Basic benchmark data set.

equivalent to the common receiver operating characteristic (ROC) curve with detection and false alarm rates.

Although the segment level classification accuracy is only slightly better than random guess [Fig. 5(a)], our OR-based fusion scheme for image level classification proves to be powerful—70% precision, 70% recall [Fig. 5(b)]. We provide a theoretical explanation of the performance gain by the fusion scheme in Section V. In Fig. 5(b), we compare our results against two random guess schemes. The first one is conducted at the image level (referred to as random guess 1), based on the population of spliced images in the data set. The second one is to first perform random guess on the segment level and then fuse these results to obtain image level decisions (random guess 2) using the same OR function. The PR curves show that the second scheme is better than the first one. Nevertheless, when compared with the experimental results, it is clear that our detector is superior to both random guess schemes.

C. Generalization to JPEG Compressed Images

In order to evaluate the performance of our algorithm under realistic scenarios where most images are JPEG compressed, we perform JPEG compression of three different quality factors ($Q = 90, 75,$ and 60) on the Basic data set and apply our SVM based splicing detector. As shown in Fig. 5(a) and (b), without retraining the classifiers, there is very little performance degradation caused by compression. Although heavy compression indeed leads to lower segment level detection accuracy [Fig. 5(a)], at image level, the recall rate only suffers a 5% drop at precision 70%. One interesting note is that at segment level, images compressed with a high quality factor ($Q = 90$) even outperform the original uncompressed data. One conjecture, though lacking empirical verification, is that high-quality compression may help remove certain image noise and achieve more accurate CRF estimation and splicing detection.

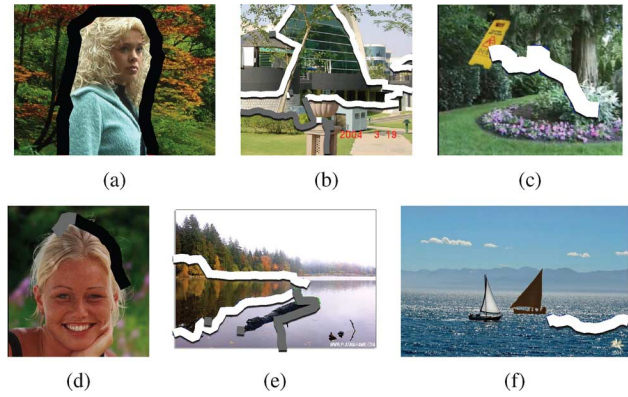


Fig. 6. Different types of spliced images (from the Advanced data set) detected by our method. Black: successfully detected spliced segments; gray: partially aligned segments detected as spliced; white: authentic segments detected as spliced.

D. Generalization to an Unseen, Challenging Data Set

When tested over the unseen Advanced data set, we observe performance decrease at segment level as anticipated: the PR curve is almost only as good as random guess [Fig. 5(c)]. Nevertheless, at image level, a precision-recall of 70% and 70% can still be obtained, comparable to that of the Basic data set [Fig. 5(d)]. Note this suggests satisfactory detection even when the trained detector is applied to new images with significantly different splicing and postprocessing operations.

In Fig. 6, we observe that spliced images with a large object, e.g., human face or body, are more likely to get precise segmentation and correct detection, even when postprocessing is present [Fig. 6(a) and (d)]. Images with objects similar to backgrounds suffer from inaccurate segmentation. However, some of the resultant partially aligned segments can still be useful, as shown in Fig. 6(b) and (e). Lastly, images with complex textures [e.g., trees in Fig. 6(c) and lake reflections in Fig. 6(f)] are prone to false alarms. In the last case, some images are detected due to segment level false alarms. To examine the effect of such “lucky” detection, we re-evaluate the performance to include only detection due to correct segment level detection and found the overall accuracy slightly degrading to 63.9%.

E. Possible Attacks to Our System

Various attack strategies exist against our system. For example, one may create a spliced image and apply another CRF so that the image appears to be generated from only one camera. Though this may weaken inconsistency cues, applying a CRF will not completely conceal the differences of the original CRFs and our method should be able to at least partially capture such inconsistency. Similarly, if the splicing is created from two images with the same CRF, we expect the boundary self fitting anomalies to still exist. Therefore, our design has a good chance to survive against these two types of attacks.

Some other tampering involves local, subtle touch-ups. In that case, if the magnitude of alteration does not exceed a certain level, it is likely to remain undetectable by our system either due to the miss in automatic segmentation or in segment level splicing detection. This is, however, beyond the scope of our work and is not the type we had originally targeted.

We offer theoretical analysis in Section V for the segment-to-image level performance gain.

V. SEGMENT-TO-IMAGE LEVEL PERFORMANCE GAIN: THEORETICAL EXPLANATION

One may question why the image level classification results appear to be much better than the segment level performance. In this section, we present theoretical validation, demonstrating how to arrive at image level results from segment level precision-recall through the simple OR fusion scheme.

Assuming the segment level false alarm rate α_S and recall β_S at a certain decision threshold τ , the image level recall β_I using our OR fusion scheme is the probability that at least one of the segments gets a score higher than the threshold

$$\begin{aligned}\beta_I &= 1 - p(d_s < \tau, \forall s \in I_S) \\ &= 1 - (1 - \alpha_S)^{n_a} (1 - \beta_S)^{n_s} (1 - \gamma_S)^{n_p}\end{aligned}\quad (13)$$

where d_s is the SVM score for segment s in a spliced image I_s , the number of authentic segments in a spliced image is denoted as n_a , the number of spliced segments n_s , and the number of partially aligned segments n_p . The quantity γ_S is the probability that the score of a partially aligned segment exceeds the threshold. Note when computing β_I , authentic images are not involved.

In the following, the uppercase letters in the subscript S and I denote whether the analysis is at the *segment* or *image* level. The lowercase letters a, s, p denote the categories of segments or images: *authentic*, *spliced*, or *partially aligned* (which is only relevant to segments but not images).

A. Conditional Probabilistic Model for OR Fusion

Equation (13) assumes every image in the data set has the same error statistics, α_S , β_S , and γ_S . However, from our experiments, we discover that images with different characteristics usually result in different detection performance.

To partially alleviate the problem, we group the images according to the distribution of segment types. We introduce a three-dimensional vector \mathbf{n} to denote the numbers of segments of different types $\{n_a, n_s, n_p\}$ in a spliced image. We collect all images with the same $\mathbf{n}_1 = \{n_{a1}, n_{s1}, n_{p1}\}$ and compute segment level $\{\alpha_{S1}, \beta_{S1}, \gamma_{S1}\}$ for this group. Below we refine the model into a conditional probabilistic formulation.

The segment level detection measures $\{\alpha_S, \beta_S, \gamma_S\}$ will be conditioned on \mathbf{n} . Since the image level recall β_I is obtained through $\{\alpha_S, \beta_S, \gamma_S\}$ and \mathbf{n} , it is also conditioned on \mathbf{n}

$$\begin{aligned}\beta_{I,\mathbf{n}} &= 1 - p(d_s < \tau, \forall s \in I_S | \mathbf{n}) \\ &= 1 - (1 - \alpha_{S,\mathbf{n}})^{n_a} (1 - \beta_{S,\mathbf{n}})^{n_s} (1 - \gamma_{S,\mathbf{n}})^{n_p}.\end{aligned}\quad (14)$$

From this formulation, we are ready to derive the image level recall, image level false alarm, and image level precision.

B. Image Level Recall, False Alarm, and Precision

The overall *image level recall* β_I is obtained by summing all conditional recalls given in (14) over all possible \mathbf{n} 's

$$\beta_I = \sum_{\mathbf{n}} \beta_{I,\mathbf{n}} p(\mathbf{n}) \quad (15)$$

where $p(\mathbf{n})$ is proportional to the number of the spliced images with distribution type \mathbf{n} .

The *Image level false alarm* α_I will be derived through authentic images. We conduct a similar conditioning procedure and obtain the conditional segment level false alarms α_{S,n_0} associated with authentic images (where n_0 denotes the number of segments within an authentic image).

The image level false alarm α_{I,n_0} is, therefore, the probability that an authentic image is classified as spliced given that it has n_0 authentic segments

$$\alpha_{I,n_0} = 1 - p(d_s < \tau, \forall s \in I | n_0) = 1 - (1 - \alpha_{S,n_0})^{n_0}. \quad (16)$$

Summing over all possible n_0 's, we get the overall image level false alarm α_I

$$\alpha_I = \sum_{n_0} \alpha_{I,n_0} p(n_0). \quad (17)$$

Having obtained the image level recall β_I and false alarm α_I in (15) and (17), the *image level precision* can be readily derived as their weighted average

$$\nu_I = \frac{\beta_I N_s}{\beta_I N_s + \alpha_I N_a} \quad (18)$$

where N_s and N_a are the numbers of spliced and authentic images, respectively.

C. Experimental Verification of Theoretical Analysis

The objective of this subsection is to verify our theoretical derivations with experimental results. The statistics from the Basic data set (Section IV) are used in this procedure. Specifically, we group the spliced images according to their \mathbf{n} 's (i.e., $\{n_a, n_s, n_p\}$) and evaluate the image level recall value $\beta_{I,\mathbf{n}}$ and segment level performance $\alpha_{S,\mathbf{n}}, \beta_{S,\mathbf{n}}, \gamma_{S,\mathbf{n}}$ for different \mathbf{n} 's and thresholds τ . From the data set, we also obtain false alarms α_{S,n_0} associated with authentic images. Probability mass functions $p(\mathbf{n})$ and $p(n_0)$ are then constructed. These values are used to predict image level performance $\{\nu_I, \beta_I\}$.

We present the comparison between predicted and actual experimental results in Fig. 7. While they may not coincide with each other exactly, our theoretical prediction offers a reasonable performance estimation and is especially useful in assessing the relative gain when fusing segment level scores to image level decisions.

VI. FEATURE SELECTION

To further discover the main contributing factor to effective splicing detection, we perform feature selection on the 20-dimensional feature vector. Re-examining (8) and (9), the original features can be categorized into different groups with distinct physical meanings:

- 1) Consistency: two-area (A, B) cross fitting

$$\begin{aligned}\mathcal{F}_{AB} &= [\mu(\mathbf{s}_{AA}), \mu(\mathbf{s}_{BB}), \mu(\mathbf{s}_{AB}), \mu(\mathbf{s}_{BA}); \\ &\quad \sigma(\mathbf{s}_{AA}), \sigma(\mathbf{s}_{BB}), \sigma(\mathbf{s}_{AB}), \sigma(\mathbf{s}_{BA}); \\ &\quad \mu(\mathbf{R}_A), \mu(\mathbf{R}_B), \Delta(\mathbf{R}_A), \Delta(\mathbf{R}_B)].\end{aligned}$$

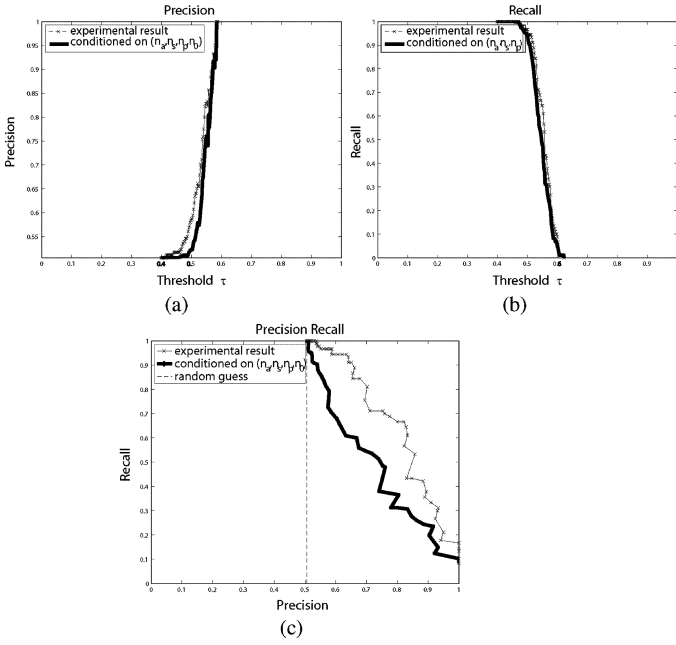


Fig. 7. Image level PR curves: predicted versus actual (a) precision, (b) recall, and (c) precision-recall curve.

2) Anomaly: self fitting of boundary segment (E)

$$\mathcal{F}_{EE} = [\mu(\mathbf{s}_{EE}), \sigma(\mathbf{s}_{EE}), \mu(\mathbf{R}_E), \Delta(\mathbf{R}_E)].$$

3) Anomaly: Self fitting of the whole area (O)

$$\mathcal{F}_{OO} = [\mu(\mathbf{s}_{OO}), \sigma(\mathbf{s}_{OO}), \mu(\mathbf{R}_O), \Delta(\mathbf{R}_O)].$$

The first group represents the consistency between two “authentic” regions, the second extracts the information from splicing boundaries, and the third captures the anomaly in the whole image. Interesting issues thus arise:

- Is two-way cross fitting between A and B sufficient for splicing detection?
- How important is the anomaly from splicing boundaries (solely from E and/or collectively from O)?
- If E and O are indeed crucial, what is the sensible way to use them?
- Can the anomaly related features be further amplified?
- Based on these tests and observations, what would be the ideal feature set for image splicing detection?

In order to answer these questions, we arrange corresponding feature sets as follows:

A) Two-way Cross Fitting

This question can be answered by the detection performance of \mathcal{F}_{AB} compared with the whole feature set.

B) Boundary Self Fitting as Standalone Features

Here we run SVM training and testing on \mathcal{F}_{EE} alone and \mathcal{F}_{OO} alone. We compare them against the results using only \mathcal{F}_{AB} to determine which factor plays the key role: two-way cross fitting from authentic regions or the anomaly created around splicing boundaries.

C) Boundary Self Fitting as Auxiliary Features

In addition to the above comparison, we run tests on $\mathcal{F}_{AB} + \mathcal{F}_{EE}$ and $\mathcal{F}_{AB} + \mathcal{F}_{OO}$ to observe the impact of \mathcal{F}_{EE} and \mathcal{F}_{OO} when combined with \mathcal{F}_{AB} . Results from

the previous and the current arrangements will reveal the optimal way of using boundary self fitting anomaly.

D) Amplified Anomaly

We further create another subset which treats areas A and B as one single source C (implicitly assuming they come from the same camera) and conduct cross fitting between this area and area E.

$$\mathcal{F}_{CE} = [\mu(\mathbf{s}_{CC}), \mu(\mathbf{s}_{EE}), \mu(\mathbf{s}_{CE}), \mu(\mathbf{s}_{EC}), \sigma(\mathbf{s}_{CC}), \sigma(\mathbf{s}_{EE}), \sigma(\mathbf{s}_{CE}), \sigma(\mathbf{s}_{EC}), \mu(\mathbf{R}_C), \mu(\mathbf{R}_E), \Delta(\mathbf{R}_C), \Delta(\mathbf{R}_E)].$$

The intuition is that when treating the content from different cameras as from the same camera, an abnormal result in CRF estimation will occur. By fitting this abnormal CRF to the abnormal (Q, R) points generated from suspected splicing area E, we can amplify the anomaly effect from two separate sources. A set \mathcal{F}_{CO} is created following the same rationale.

E) Ideal Feature Subset

Based on the findings of all previous arrangements, we arrive at a final, optimal subset of features that preserve only the crucial components to successful detection. SVM training and testing will be conducted to verify the detection power of such feature subset.

VII. FEATURE SELECTION RESULTS

In this subsection, we analyze the performance of different feature subsets described in Section VI.

A. Two-Way Cross Fitting

When excluding self fitting scores from splicing boundaries, two-way cross fitting between areas A and B pulls the classification performance down, as shown in Fig. 8(a), (e). At segment level [Fig. 8(a)], it can be as dramatic as a 15% precision decrease when recall is low (around 20%). In Fig. 8(e), the image level performance suffers even more than segment level: with recall at an acceptable 80%, precision is 10% lower, and when recall drops to as low as 30%, precision falls by 20%.

B. Boundary Self Fitting as Standalone Features

Results from the previous subsection confirm the necessity of splicing boundary information. In this subsection, we use self fitting scores from area E (or area O) as standalone features and compare the results against two-way cross fitting.

In Fig. 8(b) and (f), it is worth noting that \mathcal{F}_{EE} performs only similarly as \mathcal{F}_{AB} , demonstrating that the anomaly from splicing boundaries would not work alone. However, the anomaly included in the whole image \mathcal{F}_{OO} is comparable, or even better than the original feature set, if we choose thresholds corresponding to recall rates higher than 45%.

We also combine \mathcal{F}_{EE} and \mathcal{F}_{OO} , equivalent to collecting all the anomaly related features in the original feature set. The segment level PR curve lies between \mathcal{F}_{AB} and the original set [Fig. 8(b)]. At image level, the PR curve is very close to the original set, especially when recall is above 70% [Fig. 8(f)].

C. Boundary Self Fitting as Auxiliary Features

We now look at how anomaly related features (\mathcal{F}_{EE} and \mathcal{F}_{OO}) behave as auxiliary sets when added to the poorly per-

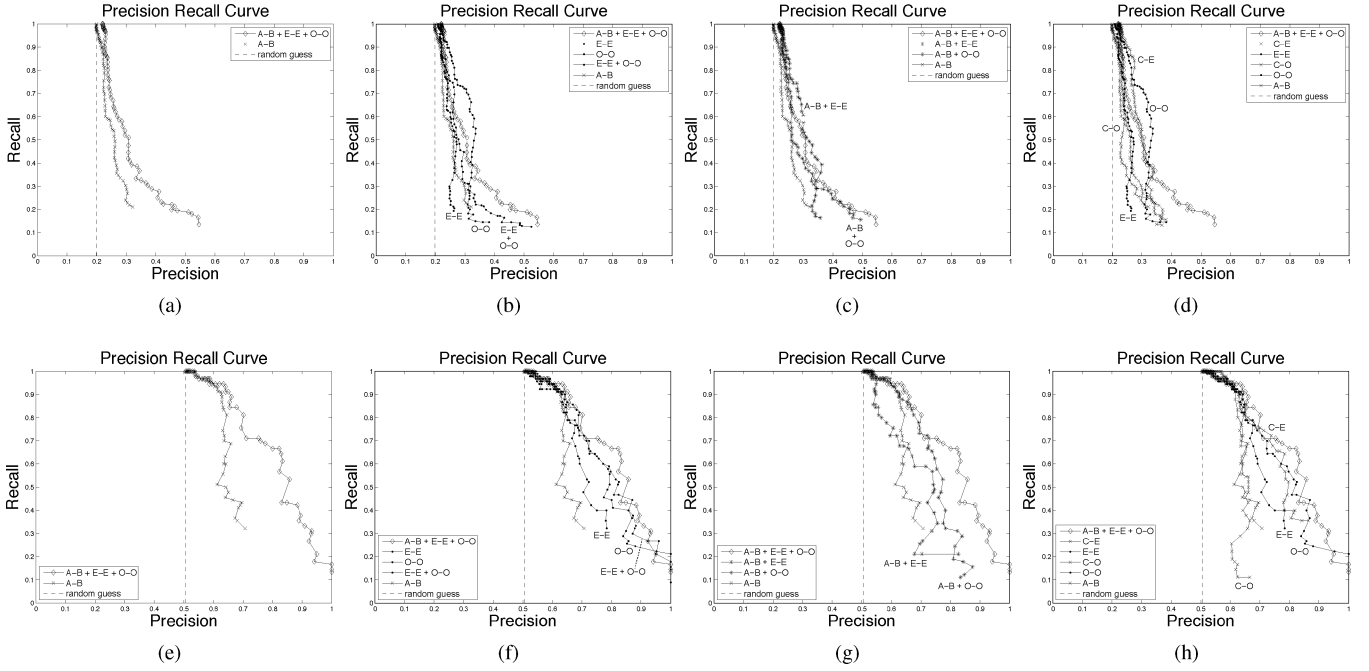


Fig. 8. (Top row: segment level; bottom row: image level) PR curves: (a), (e) two-way cross fitting; (b), (f) boundary self fitting as standalone features; (c), (g) boundary self fitting as auxiliary features; (d), (h) amplified anomaly.

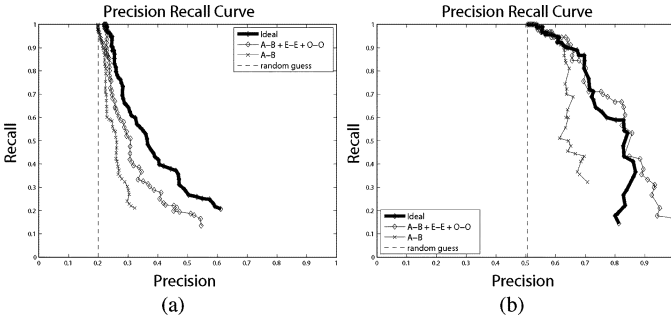


Fig. 9. PR curves: ideal feature set.

forming two-way cross fitting (\mathcal{F}_{AB}). Fig. 8(c) shows that \mathcal{F}_{EE} and \mathcal{F}_{OO} both boost the segment level performance of \mathcal{F}_{AB} , with \mathcal{F}_{EE} achieving much higher performance gain. At image level [Fig. 8(g)], \mathcal{F}_{EE} with \mathcal{F}_{AB} does not perform consistently better, while \mathcal{F}_{OO} with \mathcal{F}_{AB} performs extremely well: comparable to the whole set at recall above 70% and 5%–10% higher precision than \mathcal{F}_{EE} with \mathcal{F}_{AB} throughout.

Results show that the splicing induced anomaly is key to successful detection in all possible ways of use, whether competing with the nonanomaly related features (\mathcal{F}_{AB}), or serving as its aid when used in a collaborative way.

D. Amplified Anomaly

Previous findings inspire us to extract anomaly by treating areas A and B as one authentic area C. There is hence potential anomaly in the estimated CRF and fitting of locally planar points if they are actually from different sources. We then conduct cross fitting between this area and area E. If an image is spliced, there will be anomaly both in C and E. Compared to the

feature using \mathcal{F}_{EE} alone and \mathcal{F}_{EE} with \mathcal{F}_{AB} , in this new feature set, we amplify the significance of the anomaly. PR curves in Fig. 8(d) and (h) validate this claim.

E. Ideal Feature Set

Previous experiments have revealed the following:

- 1) Two-way cross fitting from authentic areas is not sufficient for splicing detection.
- 2) Anomaly from “hybrid” areas is crucial, especially when cross fitting between different areas.

These observations lead us to look for an ideal feature set, making use of the anomaly in the most powerful way. We take cross fitting components from \mathcal{F}_{AB} and the full \mathcal{F}_{CE} , constructing an optimal set

$$\mathcal{F}_{\text{ideal}} = [\mu(\mathbf{s}_{AB}), \mu(\mathbf{s}_{BA}), \sigma(\mathbf{s}_{AB}), \sigma(\mathbf{s}_{BA}), \mu(\mathbf{R}_A), \mu(\mathbf{R}_B), \Delta(\mathbf{R}_A), \Delta(\mathbf{R}_B), \mu(\mathbf{s}_{CC}), \mu(\mathbf{s}_{EE}), \mu(\mathbf{s}_{CE}), \mu(\mathbf{s}_{EC}), \sigma(\mathbf{s}_{CC}), \sigma(\mathbf{s}_{EE}), \sigma(\mathbf{s}_{CE}), \sigma(\mathbf{s}_{EC}), \mu(\mathbf{R}_C), \mu(\mathbf{R}_E), \Delta(\mathbf{R}_C), \Delta(\mathbf{R}_E)].$$

Capturing the contribution from both the anomaly and cross fitting, this set performs significantly better than all subsets experimented above (Fig. 9). It is even superior to the original set, especially over the range of precision from 70% to 80%.

VIII. CONCLUSION AND FUTURE WORK

We proposed an automatic spliced image detection approach based on geometry invariant CRF estimation, consistency checking, and image segmentation. The method is fully passive and automatic—neither active signature embedding nor manual input is needed. Though the accuracy in detecting specific location is modest, the segment to image level fusion not only works

well on the basic benchmark set but also gives satisfactory accuracies on a challenging realistic set. Results demonstrate that our method is promising for real-world applications. Such good performance is also verified through theoretical analysis.

With extensive feature selection, we show that the anomaly induced near the splicing boundary is crucial to successful detection, whether used as a standalone feature set or as an auxiliary set to other features. This comparative study has also provided deeper physical insight to design an optimal feature set which indeed demonstrates superior performance.

Future work may include incorporation of spatial constraints such as boundary smoothness, continuity, and alignment with image scene context to prune classification errors and refine the localization accuracy of specific spliced areas.

ACKNOWLEDGMENT

The authors would like to thank T. Ng for valuable discussions and codes and J. He for assistance with the Advanced data set from MSRA.

REFERENCES

- [1] A. Popescu and H. Farid, "Exposing digital forgeries in color filter array interpolated images," *IEEE Trans. Signal Process.*, vol. 53, no. 10, pt. 2, pp. 3948–3959, Oct. 2005.
- [2] A. Swaminathan, M. Wu, and K. Liu, "Nonintrusive component forensics of visual sensors using output images," *IEEE Trans. Inf. Forensics Security*, vol. 2, no. 1, pp. 91–106, Mar. 2007.
- [3] A. Swaminathan, M. Wu, and K. Liu, "Digital image forensics via intrinsic fingerprints," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 101–117, Mar. 2008.
- [4] J. Lukas, J. Fridrich, and M. Goljan, "Digital camera identification from sensor pattern noise," *IEEE Trans. Inf. Forensics Security*, vol. 1, no. 2, pp. 205–214, Jun. 2006.
- [5] Y. Tsin, V. Ramesh, and T. Kanade, "Statistical calibration of the CCD imaging process," in *Proc. IEEE Int. Conf. Computer Vision*, 2001, pp. 480–487.
- [6] H. Farid, "Blind inverse gamma correction," *IEEE Trans. Image Process.*, vol. 10, no. 10, pp. 1428–1433, Oct. 2001.
- [7] S. Lin, J. Gu, S. Yamazaki, and H.-Y. Shum, "Radiometric calibration from a single image," in *Proc. Int. Conf. Computer Vision and Pattern Recognition*, Washington, DC, 2004, vol. 2, pp. 938–945.
- [8] T.-T. Ng, S.-F. Chang, and M.-P. Tsui, "Using geometry invariants for camera response function estimation," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, Minneapolis, MN, Jun. 2007, pp. 1–8.
- [9] S. Lin and L. Zhang, "Determining the radiometric response function from a single grayscale image," in *Proc. Int. Conf. Computer Vision and Pattern Recognition*, 2005, pp. 66–73.
- [10] M. Kharrazi, H. T. Sencar, and N. D. Memon, "Blind source camera identification," in *Proc. Int. Conf. Image Processing*, 2004, pp. 709–712.
- [11] M. Chen, J. Fridrich, M. Goljan, and J. Lukas, "Determining image origin and integrity using sensor noise," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 74–90, Mar. 2008.
- [12] Z. Lin, R. Wang, X. Tang, and H.-Y. Shum, "Detecting doctored images using camera response normality and consistency," in *Proc. Int. Conf. Computer Vision and Pattern Recognition*, San Diego, CA, 2005, vol. 1, pp. 1087–1092.
- [13] Y.-F. Hsu and S.-F. Chang, "Detecting image splicing using geometry invariants and camera characteristics consistency," in *Proc. Int. Conf. Multimedia and Expo*, Toronto, ON, Canada, 2006, pp. 549–552.
- [14] J. He, Z. Lin, L. Wang, and X. Tang, "Detecting doctored JPEG images via DCT coefficient analysis," in *Proc. 9th Eur. Conf. Computer Vision*, Graz, Austria, 2006, vol. 3954, pp. 423–435.
- [15] S. Bayram, I. Avcibas, B. Sankur, and N. Memon, "Image manipulation detection," *J. Electron. Imaging*, vol. 15, no. 4, p. 041102, Oct./Dec. 2006.
- [16] S. Mann, "Comparametric equations with practical applications in-quantigraphic image processing," *IEEE Trans. Image Process.*, vol. 9, no. 8, pp. 1389–1406, Aug. 2000.
- [17] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *Proc. 24th Annu. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH'97)*, New York, NY, 1997, pp. 369–378.
- [18] M. Grossberg and S. Nayar, "What can be known about the radiometric response from images?," in *Proc. 7th Eur. Conf. Computer Vision*, 2002, pp. 189–205.
- [19] M. Grossberg and S. Nayar, "What is the space of camera response functions?," in *Proc. Int. Conf. Computer Vision and Pattern Recognition*, Madison, WI, 2003, vol. 2, pp. 602–609.
- [20] Y.-F. Hsu, "Image Tampering Detection for Forensics Applications," Ph.D. Dissertation, Graduate School of Arts and Sciences, Columbia Univ., New York, 2009.
- [21] Y.-F. Hsu and S.-F. Chang, "Statistical fusion of multiple cues for image tampering detection," in *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, 2008, pp. 1386–1390.
- [22] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 888–905, Aug. 2000.
- [23] C. Yang, R. Duraiswami, N. Gumerov, and L. Davis, "Improved fast gauss transform and efficient kernel density estimation," in *Proc. Ninth IEEE Int. Conf. Computer Vision*, 2003, pp. 664–671.
- [24] Y.-F. Hsu and S.-F. Chang, "Image splicing detection using camera response function consistency and automatic segmentation," in *Proc. Int. Conf. Multimedia and Expo*, Beijing, China, 2007, pp. 28–31.



Yu-Feng Hsu received the B.S. degree in electrical engineering from National Taiwan University and the M.S. degree from Carnegie Mellon University. She received the Ph.D. degree in electrical engineering from Columbia University, New York, in 2009.

She is currently a researcher at the Information and Communications Research Laboratories, Industrial Technology Research Institute, Taiwan. Her research interest lies along the line of digital image forensics and video content analysis, including watermarking, biometrics, and photograph tampering detection

using device characteristics and statistical machine learning. Her current focus is computer vision algorithms for smart video surveillance systems.



Shih-Fu Chang (S'89-M'90-SM'01-F'04) is Professor and Chairman of Electrical Engineering and Director of the Digital Video and Multimedia Laboratory, Columbia University, New York. He has made significant contributions in multimedia search, media forensics, video adaptation, and international standards for multimedia indexing. He has worked in different advising/consulting capacities for many media technology companies and research institutions.

Dr. Chang has been recognized with several awards, including the IEEE Kiyo Tomiyasu Award, the Navy ONR Young Investigator Award, the IBM Faculty Award, the ACM Recognition of Service Award, and the NSF CAREER Award. He and his students have received several Best Paper and Best Student Paper Awards from IEEE, ACM, and SPIE. He was Editor-in-Chief for *IEEE Signal Processing Magazine* from 2006 to 2008.