Digital Image Authentication from JPEG Headers

Eric Kee, Micah K. Johnson and Hany Farid

***Abstract*— It is often desirable to determine if an image has been modified in any way from its original recording. The JPEG format affords engineers many implementation trade-offs which give rise to widely varying JPEG headers. We exploit these variations for image authentication. A camera signature is extracted from a JPEG image consisting of information about quantization tables, Huffman codes, thumbnails, and EXIF format. We show that this signature is highly distinct across 1.3 million images spanning 773 different cameras and cellphones. Specifically, 62% of images have a signature that is unique to a single camera, 80% of images have a signature that is shared by three or fewer cameras, and 99% of images have a signature that is unique to a single manufacturer. The signature of Adobe Photoshop is also shown to be unique relative to all 773 cameras. These signatures are simple to extract and offer an efficient method to establish the authenticity of a digital image.**

I. Introduction

Digital images are now routinely introduced as evidence into a court of law. It has, therefore, become critical to verify the integrity of this digital evidence. Digital forensic techniques have been developed to detect various traces of tampering: region duplication [1], [2], [3], [4]; resampling [5], [6]; color filter array artifacts [7], [8]; inconsistencies in camera response function [9]; inconsistencies in lighting and shadows [10], [11], [12], [13]; inconsistencies in chromatic aberrations [14], [15]; inconsistencies in sensor noise [16], [17]; and inconsistencies in statistical features [18]. See [19] for a general survey. However, relatively benign modifications either cannot be detected by these techniques, or render these techniques ineffective.

It is often desirable to determine if a digital image has been altered in any way from the time of its recording, including manipulations as simple as cropping. Previous work in detecting double JPEG compression holds some promise to detect any form of image manipulation [20], [21], [22], [23], [24]. These techniques require fairly sophisticated models and analysis schemes, can be vulnerable to simple counter-measures such as additive noise or down-sampling, and can be computationally intensive. In contrast, it has been previously shown that EXIF headers [25] and JPEG quantization tables [26], [27], [28], [29] used by cameras and software manufacturers are somewhat distinct, and can therefore be used to determine if an image has been altered from its original recording. Building on this earlier work, we describe how various aspects of the JPEG format can be used for authentication. Unlike previous work, this approach considers several features of the JPEG format not previously considered, namely properties of the run-length encoding employed by the JPEG standard, and aspects of the EXIF header format. This analysis is validated on over 1.3 million images spanning 33 different camera manufacturers and 773 different camera and cellphone models.

E. Kee (erickee@cs.dartmouth.edu) and H. Farid (farid@cs.dartmouth.edu) are with the Department of Computer Science at Dartmouth College, 6211 Sudikoff Lab, Hanover NH 03755. M.K. Johnson (kimo@csail.mit.edu) is with the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology, 32 Vassar Street, 32-D462 Cambridge, MA 02139.

II. Methods

The JPEG file format has emerged as a universal image standard employed by nearly all commercial digital cameras [30], [31]. As such, we consider the details of this encoding scheme, and how these details vary among cameras of different make, model, resolution, and quality. Cameras often support multiple resolutions and qualities, each of which yield images with different JPEG compression parameters. We will, therefore, explore the JPEG parameters used by cameras of different make and model, and by each camera under different resolution and quality settings.

We begin by describing the JPEG compression standard and file format. Given a three channel color image (RGB), compression proceeds as follows. An image is first transformed from RGB into luminance/chrominance space (YCbCr). The two chrominance channels (CbCr) are typically subsampled by a factor of two relative to the luminance channel (Y). Each channel is then partitioned into 8×8 pixel blocks. These values are converted from unsigned to signed integers (e.g., from $[0,255]$ to $[-128,127]$). Each block, $f_c(\cdot)$, is converted to frequency space, $F_c(\cdot)$, using a 2-D discrete cosine transform (DCT):

$$F_c(u,v) = \alpha_{u,v} \sum_{x=0}^{7}\sum_{y=0}^{7} f_c(x,y) \cos\left(\frac{(2x+1)u\pi}{16}\right) \cos\left(\frac{(2y+1)v\pi}{16}\right), \quad (1)$$

where c denotes a specific image channel, $\alpha_{u,v}$ is a normalizing scale factor, and $f_c(\cdot)$ is the underlying pixel values. Depending on the specific frequency u,v and channel c, each DCT coefficient, $F_c(\cdot)$, is quantized by an amount $q_c(\cdot)$:

$$\hat{F}_c(u,v) \quad = \quad \text{round}\left(\frac{F_c(u,v)}{q_c(u,v)}\right). \quad (2)$$

This stage is the primary source of data reduction and information loss.

With some variations, the above sequence of steps is employed by all JPEG encoders. The primary source of variation in these encoders is the choice of quantization values $q_c(\cdot)$. The quantization is specified as a set of three 8×8 tables associated with each frequency and image channel (YCbCr). For low compression rates, the values in these tables tend towards 1, and increase for higher compression rates. Shown in top portion of Fig. 2, for example, are the quantization tables employed by a Canon and Nikon digital camera. Shown,

from top to bottom, are the tables for the luminance and chrominance channels. As is typical, the quantization for the luminance channel is less than for the two chrominance channels, and the quantization is less for the lower frequency components (the frequency in each table is specified in standard zig-zag order with the lowest frequency in the top left corner, and the highest frequency in the bottom right corner).

After quantization, the DCT coefficients are subjected to entropy encoding (typically Huffman coding). Huffman coding is a variable-length encoding scheme that encodes frequently occurring values with shorter codes, and less frequently occurring values with longer codes. This lossless compression scheme exploits the fact that the quantization of DCT coefficients yields many zero coefficients, which can in turn be efficiently encoded. Motivated by the fact that the statistics of the DC and AC DCT coefficients are different[1], the JPEG standard allows for different Huffman codes for the DC and AC coefficients. This entropy encoding is applied separately to each YCbCr channel, employing separate codes for each channel.

The JPEG standard does not enforce any specific quantization table or Huffman code. Camera and software engineers are therefore free to balance compression and quality to their own needs and tastes. The specific quantization tables and Huffman codes needed to decode a JPEG file are embedded into the JPEG header. In the following sections we describe how the JPEG quantization table and Huffman codes along with other data extracted from the JPEG header form a distinct camera signature which can be used for authentication.

A. Image Parameters

The first three components of our camera signature are the image dimensions, quantization table, and Huffman code. The image dimensions are used to distinguish between cameras with different sensor resolution. In order to compensate for landscape and portrait images, the image dimensions are specified as the minimum dimension followed by the maximum dimension. The set of three 8×8 quantization tables are specified as a one dimensional array of 192 values: each channel's table is specified in column-order, and the three tables are specified in the order of luminance (Y), chrominance (Cb) and chrominance (Cr).

The Huffman code is specified as six sets of 15 values corresponding to the number of codes of length 1, $2 \ldots 15$: each of three channels requires two codes, one for the DC coefficients and one for the AC coefficients. This representation eschews the actual code for a more compact representation that distinguishes codes based on the distribution of code lengths.

In theory the chrominance channels, Cb and Cr, can employ different quantization values and Huffman codes. In all of the cameras analyzed, however, we have found that the chrominance channels are encoded with the same parameters.

In total, we extract 284 values from the full resolution image: 2 image dimensions, 192 quantization values, and 90 Huffman codes.

[1]The DC term refers to the $(0,0)$ frequency in the top left corner of each quantization table. The AC terms refer to the remaining frequencies.

B. Thumbnail Parameters

A thumbnail version of the full resolution image is often embedded in the JPEG header. The next three components of our camera signature are extracted from this thumbnail image. A thumbnail is typically no larger in size than a few hundred square pixels, and is created by cropping, filtering and down-sampling the full-resolution image. The thumbnail is then typically compressed and stored in the header as a JPEG image. As such, we can extract the same components from the thumbnail as from the full resolution image described in the previous section. As with the full resolution image, we have found that the chrominance channels are encoded with the same parameters.

Some camera manufacturers do not create a thumbnail image, or do not encode them as a JPEG image. In such cases, we simply assign a value of zero to all of the thumbnail parameters. Rather than being a limitation, we consider the lack of a thumbnail as a characteristic property of a camera.

In total, we extract 284 values from the thumbnail image: 2 thumbnail dimensions, 192 quantization values, and 90 Huffman codes.

C. EXIF Metadata Parameters

The final component of our camera signature is extracted from an image's EXIF metadata [32]. The metadata, found in the JPEG header, stores a variety of information about the camera and image. According to the EXIF standard, there are five main image file directories (IFDs) into which the metadata is organized: (1) Primary; (2) Exif; (3) Interoperability; (4) Thumbnail; and (5) GPS. Camera manufacturers are free to embed any (or no) information into each IFD. We extract a compact representation of their choice by counting the number of entries in each of these five IFDs.

Because the EXIF standard allows for the creation of additional IFDs, we also count the total number of any additional IFDs, and the total number of entries in each of these. Some camera manufacturers customize their metadata in ways that do not conform to the EXIF standard. These customizations yield errors when parsing the metadata. We consider these errors to be a feature of camera design and therefore count the total number of parser errors, as specified by the EXIF standard.

In total, we extract 8 values from the metadata: 5 entry counts from the standard IFDs, 1 for the number of additional IFDs, 1 for the number of entries in these additional IFDs, and 1 for the number of parser errors.

D. Image Authentication

As described in the previous sections, we extract 284 header values from the full resolution image, a similar 284 header values from the thumbnail image, and another 8 from the EXIF metadata, for a total of 576 values. These 576 values form the signature by which images will be authenticated. Specifically, the signature and camera make and model are extracted from the EXIF metadata and compared against authentic image signatures extracted from the same camera make and model.

To the extent that photo-editing software will employ JPEG parameters that are distinct from the camera's, any manipulation will alter the original signature, and can therefore be detected. Specifically, photo alteration is detected by extracting the signature from an image and comparing it to a database of known authentic camera signatures. Any matching camera make and model can then be compared to the make and model specified in the image's EXIF metadata. Any mismatch is strong evidence of some form of tampering.

For ballistic purposes, an extracted signature can be compared against authentic signatures to determine which, if any, cameras have matching signatures. This application may seem unnecessary since the camera make and model are specified in the EXIF metadata. However, an image's EXIF metadata can be relatively easily edited to alter the camera make and model, so the signature is a more reliable determinant of a camera's source.

Given that an image's EXIF metadata can be easily edited, it may seem peculiar to rely on it for forensic and ballistic purposes. It should be noted that modifying the content of any existing EXIF field will not affect the extracted EXIF counts, and hence will not affect the extracted signature. If however, the camera make and model fields are changed in an attempt to conceal the camera source, then this can be detected when the extracted signature is found to be inconsistent with the make and model. Additionally, if the metadata is accidentally or intentionally deleted, then the remaining portion of the signature can still be used for authentication. If, on the other hand, the number of EXIF fields are increased or decreased, then the EXIF count will be inconsistent with the expected signature. Although an image's EXIF metadata can be edited, it still provides useful information for forensic and ballistic analysis.

The usefulness of the camera signature is only as good as our ability to extract signatures from a multitude of cameras of different make, model, quality and resolution settings. Therefore we describe next the construction of a large database of images collected from on-line sources.

E. Image Database

Approximately 10 million images were downloaded from the popular photo-sharing website *Flickr.com*. Since we are interested in extracting original camera signatures, it was necessary to eliminate any images that had been edited or altered by photo-editing software. To begin, only images tagged as "original" by *Flickr* were downloaded. The images were then subjected to a series of filtering stages:

1) Images that were not 3-channel color JPEG images were eliminated.
2) Duplicate images, determined by comparing MD5 hashes, were eliminated.
3) Images with no metadata or reduced metadata were eliminated.
4) Images with inconsistent metadata "modification" and "original" dates were eliminated
5) Images with a metadata "software" tag, introduced or modified by a photo-editing software, were eliminated. The most common modifications were found to contain one of the following keywords:

   `adobe photoshop | aperture borderfx | ashampoo photo commander | bibble | capture nx | capture one | coachware | copiks photomapper | digikam | digital photo pro | gimp | www.idimager.com | imagenomic noiseware | imageready | kipi | microsoft | paint.net | paint shop pro photo | photoscape | photowatermark | picasa | picnik | quicktime | watermark`
6) Under certain conditions, *Flickr* resizes images to a maximum dimension of 640, 800, 1024, 1280, 1600, or 2048. When doing so, we found that *Flickr* frequently employs one of two image quantization tables (Appendix I). Images were eliminated if they were of any of these sizes and employed either of these quantization tables.
7) Images whose resolution is not native to the camera make and model were eliminated. The native resolutions of $1{,}578$ different cameras were gathered from the website *dpreview.com* using Amazon's crowd-sourcing tool Mechanical Turk. Users on Mechanical Turk were given a camera make and model and a link to *dpreview.com*. They were asked to retrieve the camera's native resolutions from the manufacturer's announcement for the camera or from the camera specification page. The data for each camera was considered valid after three independent users entered the same list of image sizes.

These filters reduced the original 10 million images to 2.2 million images.

The camera make, model, and signature were then extracted from each of these image's metadata. In order to further eliminate possible edited or altered images, only those entries with 25 or more images having the same paired make, model, and signature were retained. This yielded approximately 1.3 million images. These images span $9{,}163$ different distinct pairings of camera make, model, and signature and represent 33 different camera manufacturers and 773 different camera and cellphone models. We refer to a pairing of camera make, model, and signature as a camera configuration. Because a camera can record in multiple resolutions and qualities, it is often the case that the same make and model camera can yield many different signatures. In our case, there is an average of 12 different signatures for each camera make and model (the ratio of configurations to unique camera models is $9{,}163/773 = 11.85$).

It is from these 1.3 million images and $9{,}163$ camera configurations that the distinctness of the camera signature was analyzed. Note that any search of this data will be computationally efficient as it will only need to be performed on the $9{,}163$ configurations, and not the much larger 1.3 million images.

Finally, we note that many manufacturers employ different model names for the same camera. For example, the Canon Digital IXUS is sold under the name Canon Powershot in the United States and Canada, and IXY Digital in Japan. These

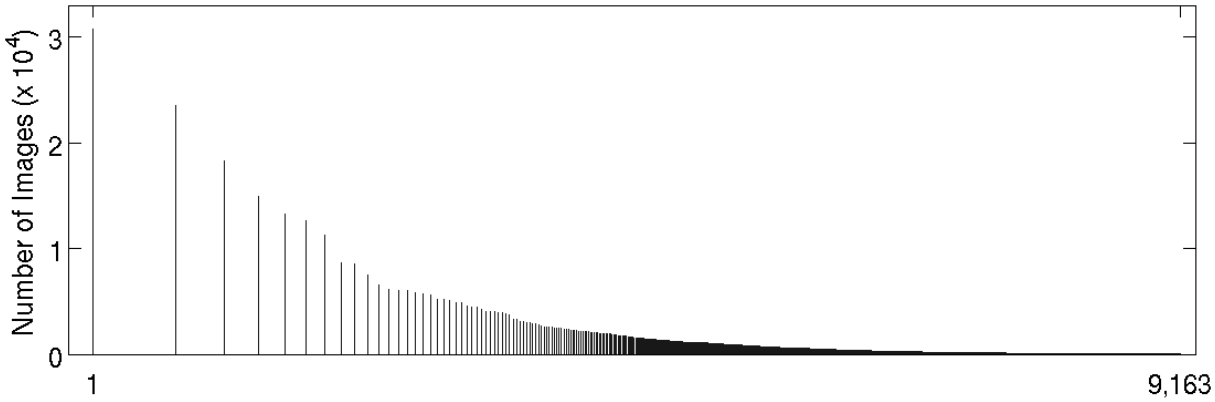

Fig. 1. Shown are, in sorted order, the number of images collected for each of 9,163 camera configurations (on a logarithmic horizontal axis).

synonymous names are taken into consideration when collating the images.

III. Results

Shown in Fig. 1 are the number of images, in sorted order, collected for each of 9,163 camera configurations. The maximum image count (30,760) is for the Canon EOS Digital Rebel XTi (also known as Canon EOS 400D Digital or EOS Kiss Digital X), and the mean and median number of images per camera configuration is 147.1 and 48, respectively. The remaining configurations span cameras and cellphones from Apple, Asahi, Canon, Casio, Fuji, Hewlett-Packard, HTC, JVC, Kodak, Konica, Leica, LG, Minolta, Motorola, Nikon, Nokia, Olympus, Panasonic, Pentax, Ricoh, Research In Motion (Blackberry), Samsung, Sanyo, Sigma, Sony, Toshiba, Vivitar, and more.

Shown in Fig. 2 are the complete signatures for a Canon EOS Rebel XTi and a Nikon D40. From top to bottom are the image dimensions, image quantization tables, image Huffman codes, thumbnail dimensions, thumbnail quantization tables, thumbnail Huffman codes, and EXIF counts.

The signatures from each of the 9,163 camera configurations were compared to determine their distinctiveness. Specifically, we determine the distinctiveness across cameras of different make and model and relative to photo-editing software. To begin, all cameras with the same signature were placed into an equivalence class. An equivalence of size of size n means that n cameras of different make and model share the same signature. An equivalence class of size greater than 1 means that there is an ambiguity in identifying the camera make and model. We would like to maximize the number of camera configurations in an equivalence class of size 1 and minimize the largest equivalence class size.

Shown in Fig. 3 and summarized in Fig. 4 are the distribution of equivalence class sizes for the signature in its whole and in its parts.

Shown in Fig. 3(a) is the distribution of equivalence class sizes for the 284-valued signature based only on the full resolution image parameters. 12.9% of the camera configurations are in an equivalence class of size one (i.e., are unique), 7.9% are in an equivalence class of size two, 6.2% are in an equivalence class of size three, 27.0% are in an equivalence class of size three or less, and the largest equivalence class is of size 185, with 2.7% of the cameras configurations.

Shown in Fig. 3(b) is the distribution of equivalence class sizes for the 284-valued signature based only on the thumbnail image parameters. 1.1% of the camera configurations are in an equivalence class of size one, 1.1% are in an equivalence class of size two, 1.0% are in an equivalence class of size three, 3.2% are in an equivalence class of size three or less, and the largest equivalence class is of size 960, with 14.1% of the camera configurations.

Shown in Fig. 3(c) is the distribution of equivalence class sizes for the 8-valued signature based only on the EXIF metadata parameters. 8.8% of the camera configurations are in an equivalence class of size one, 5.4% are in an equivalence class of size two, 4.2% are in an equivalence class of size three, 18.4% are in an equivalence class of size three or less, and the largest equivalence class is of size 188, with 2.8% of the camera configurations.

Shown in Fig. 3(d) is the distribution of equivalence class sizes for the 568-valued signature based on the full resolution image and thumbnail image parameters. 24.9% of the camera configurations are in an equivalence class of size one. 15.3% are in an equivalence class of size two, 11.3% are in an equivalence class of size three, 51.5% are in an equivalence class of size three or less, and the largest equivalence class is of size 91, with 1.3% of the camera configurations.

Shown in Fig. 3(e) is the distribution of equivalence class sizes for the complete 576-valued signature. 69.1% of the camera configurations are in an equivalence class of size one. 12.8% are in an equivalence class of size two, 5.7% are in an equivalence class of size three, 87.6% are in an equivalence class of size three or less, and the largest equivalence class is of size 14, with 0.2% of the camera configurations. Because the distribution of cameras is not uniform, it is also useful to consider the likelihood of an image, as opposed to camera configuration, being in an equivalence class of size n. With respect to the complete signature, 62.4% of images are in an equivalence class of size one (i.e., are unique), 10.5% of images are in an equivalence class of size two, 7.5% of images are in an equivalence class of size three, and 80.4% of images

	Canon EOS Rebel XTi	**Nikon D40**
image dimensions	2592 × 3888	2000 × 3008
image quantization table (Y)	1 1 1 1 1 2 3 3 1 1 1 1 1 3 3 3 1 1 1 1 2 3 3 3 1 1 1 1 3 4 4 3 1 1 2 3 3 5 5 4 1 2 3 3 4 5 6 5 2 3 4 4 5 6 6 5 4 5 5 5 6 5 5 5	1 1 1 1 1 1 1 2 1 1 1 1 1 2 2 2 1 1 1 1 1 2 2 2 1 1 1 1 1 2 2 2 1 1 1 2 2 3 3 2 1 1 2 2 2 3 3 3 1 2 2 2 3 3 3 3 2 3 3 3 3 3 3 3
image quantization table (Cb)	1 1 1 2 5 5 5 5 1 1 1 3 5 5 5 5 1 1 3 5 5 5 5 5 2 3 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5	1 1 1 1 3 3 3 3 1 1 1 2 3 3 3 3 1 1 2 3 3 3 3 3 1 2 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3
image quantization table (Cr)	1 1 1 2 5 5 5 5 1 1 1 3 5 5 5 5 1 1 3 5 5 5 5 5 2 3 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5	1 1 1 1 3 3 3 3 1 1 1 2 3 3 3 3 1 1 2 3 3 3 3 3 1 2 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3 3
image Huffman code DC (Y)	1 5 1 1 1 1 1 1 0 0 0 0 0 0 0 0	1 5 1 1 1 1 1 1 0 0 0 0 0 0 0 0
image Huffman code DC (Cb)	3 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0	3 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0
image Huffman code DC (Cr)	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
image Huffman code AC (Y)	2 1 3 3 2 4 3 5 5 4 4 0 0 1 125	2 1 3 3 2 4 3 5 5 4 4 0 0 1 125
image Huffman code AC (Cb)	2 1 2 4 4 3 4 7 5 4 4 0 1 2 119	2 1 2 4 4 3 4 7 5 4 4 0 1 2 119
image Huffman code AC (Cr)	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
thumbnail dimensions	120 × 160	120 × 160
thumbnail quantization table (Y)	3 2 2 3 5 8 10 12 2 2 3 4 5 11 11 13 3 2 3 5 8 11 13 11 3 3 4 6 10 17 15 12 3 4 7 11 13 21 20 15 5 7 10 12 15 20 21 17 9 12 15 17 20 23 23 19 14 17 18 19 21 19 20 19	1 1 1 1 1 2 3 4 1 1 1 1 2 3 4 3 1 1 1 1 2 3 4 3 1 1 1 2 3 5 5 4 1 1 2 3 4 6 6 5 1 2 3 4 5 6 7 5 3 4 5 5 6 7 7 6 4 5 6 6 7 6 6 6
thumbnail quantization table (Cb)	3 3 5 9 19 19 19 19 3 4 5 13 19 19 19 19 5 5 11 19 19 19 19 19 9 13 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19	1 1 1 3 6 6 6 6 1 1 2 4 6 6 6 6 1 2 3 6 6 6 6 6 3 4 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
thumbnail quantization table (Cr)	3 3 5 9 19 19 19 19 3 4 5 13 19 19 19 19 5 5 11 19 19 19 19 19 9 13 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19 19	1 1 1 3 6 6 6 6 1 1 2 4 6 6 6 6 1 2 3 6 6 6 6 6 3 4 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
thumbnail Huffman code DC (Y)	1 5 1 1 1 1 1 1 0 0 0 0 0 0 0 0	1 5 1 1 1 1 1 1 0 0 0 0 0 0 0 0
thumbnail Huffman code DC (Cb)	3 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0	3 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0
thumbnail Huffman code DC (Cr)	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
thumbnail Huffman code AC (Y)	2 1 3 3 2 4 3 5 5 4 4 0 0 1 125	2 1 3 3 2 4 3 5 5 4 4 0 0 1 125
thumbnail Huffman code AC (Cb)	2 1 2 4 4 3 4 7 5 4 4 0 1 2 119	2 1 2 4 4 3 4 7 5 4 4 0 1 2 119
thumbnail Huffman code AC (Cr)	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
EXIF count	9 28 2 6 0 7 162 0	10 40 2 7 0 7 94 0

Fig. 2. Camera signatures for a Canon EOS Rebel XTi and Nikon D40 highlighting the differences (image dimensions, image quantization tables, thumbnail quantization tables and EXIF counts) and similarities (thumbnail dimensions and image and thumbnail Huffman codes).

(a) image

(b) thumbnail

(c) EXIF

(d) image+thumbnail

(e) all

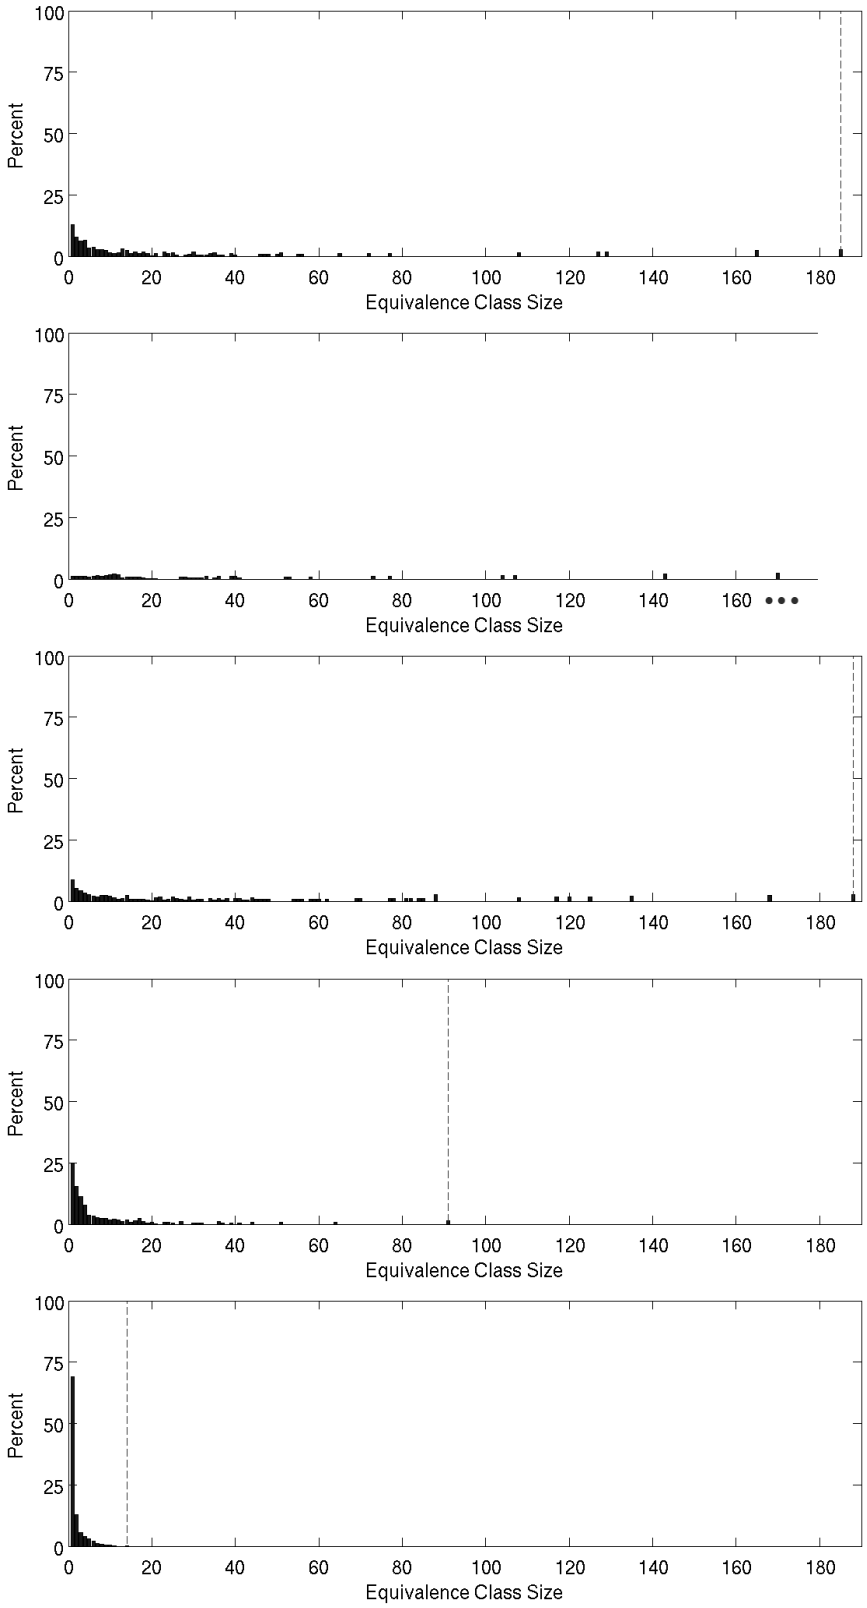

Fig. 3. Shown are the distributions of equivalence class sizes based on the distinctiveness of the signature (an equivalence class of size n means that n cameras of different make and model share the same signature). Panels (a) through (e) correspond to, from top to bottom, signatures consisting of image parameters only, thumbnail parameters only, EXIF parameters only, image and thumbnail parameters, and image, thumbnail and EXIF parameters. The dashed vertical line in each panel corresponds to the maximum equivalence class size. The horizontal axis for panel (b) was truncated (the maximum equivalence class size is 960). See also Fig. 4.

	Equivalence Class Size						
	1	2	3	4	5	median	max
image	12.9%	7.9%	6.2%	6.6%	3.4%	11	185
thumbnail	1.1%	1.1%	1.0%	1.1%	0.7%	694	960
EXIF	8.8%	5.4%	4.2%	3.2%	2.6%	25	188
image + thumbnail	24.9%	15.3%	11.3%	7.9%	3.7%	3	91
all	69.1%	12.8%	5.7%	4.0%	2.9%	1	14

Fig. 4. Shown are the percentage of camera configurations with an equivalence class size of $1 \ldots 5$, and the median and maximum equivalence class size. Each row corresponds to different subsets of the complete signature. See also Fig. 3.

are in an equivalence class of size three or less.

Individually, the image, thumbnail and EXIF parameters are not particularly distinct, Fig. 3(a-c), but when combined, they provide a highly distinct signature, Fig. 3(e). This suggests that the choice of parameters are not highly correlated, and hence their combination improves overall distinctiveness. Although the thumbnail parameters are the least distinct, their addition to the overall signature is significant as can be seen by comparing Fig. 3(a) and (d), and the second and fourth rows of Fig. 4.

Shown in Fig. 5(a) are the relative contributions of the individual parameters. The horizontal bars denote the number of unique values for each feature. The EXIF count is most unique, followed by image dimensions, image quantization table (Y channel) and then thumbnail quantization table (Y channel). The least distinct parameters are the Huffman tables and the thumbnail dimensions. Shown in Fig. 5(b)-(e) is this same analysis for all Canon, Sony, Nikon, and Olympus cameras: the Canon cameras are considerably more consistent than other manufacturers. This consistency will yield to ambiguities in identifying a camera based on its signature. As shown next, these ambiguities are almost always manufacturer specific.

An equivalence class of size greater than 1 implies a non-distinct signature, and hence an ambiguity in identifying the camera make and model. We next consider the scope of this ambiguity. Shown in Fig. 6 are the cameras' make and model that are in the same equivalence class for the full 576-valued signature. Shown in the last two rows, for example, are the cameras in the largest equivalence classes of size of 11 and 14. Note that all of these cameras are variants of the Sony DSC series. Each of the four equivalence classes of size 10 consist exclusively of the Sony DSC series. Note that, because of the multiple resolution and quality settings, the same camera can appear in multiple equivalence classes. Similarly, each of the seven equivalence classes of size 8 and each of the eleven equivalence classes of size 7 consist exclusively of either the Sony DSC series, the Canon Powershot series, or the Research In Motion (RIM) BlackBerry. Although not shown in Fig. 6, this pattern continues for nearly all equivalence classes of size greater than 1. With only a few exceptions, each equivalence class consists of cameras from the same manufacturer and series. The only exceptions are the following four equivalence classes of size two:

- Casio EX-Z60 | Canon Powershot SX120 IS
- Nikon Coolpix P90 | Panasonic DMC-FZ18
- Panasonic DMC-TZ5 | Nikon Coolpix S52
- Panasonic DMC-ZS7 | Nikon Coolpix S630

In summary, although there is an ambiguity in some of the signatures, the signature still significantly constrains the identity of the camera make and model.

Lastly, the signature from Adobe Photoshop (versions 3, 4, 7, CS, CS2, CS3, CS4, CS5 at all qualities) were compared to the $9{,}163$ camera signatures. In this case, only the image and thumbnail quantization tables and Huffman codes were used for comparison. No overlap was found between any Photoshop version/quality and camera manufacturer. As such, the Photoshop signatures, each residing in an equivalence class of size 1, are unique. This means that any photo-editing with Photoshop can be easily and unambiguously detected.

IV. Discussion

We have shown that cameras produce distinct JPEG headers that facilitate both forensic and ballistic analysis. This analysis does not differentiate between benign and nefarious modifications. While this is a stringent criteria, it is useful in certain arenas. This forensic analysis can be useful in a legal setting, for example, where it is important to determine if evidence has been altered in any way.

As with any forensic technique, it is important to consider counter-measures. In our case, a determined forger could conceal their traces of tampering by extracting the signature of a camera, modifying the image, and then re-saving the image with the appropriate EXIF format and all of the appropriate parameters: image size, image quantization table, image Huffman code, thumbnail size, thumbnail quantization table, and thumbnail Huffman code. While this is certainly possible, it is currently beyond the scope of popular photo-editing software. Our analysis is also vulnerable to a standard re-broadcast attack in which a digital image is manipulated, printed, and re-photographed.

Because an image's EXIF metadata can be easily edited, this analysis can provide a more reliable method to determine camera make and model. This information can be useful in other forensic analyses. For example, techniques for device authentication (e.g., [16]) compare an image of unknown origin to a database of known cameras. For large databases, this analysis can be computationally demanding [33], [34], [35]. Reliable determination of camera make and model can reduce this complexity by focusing only on the relevant camera(s) in the database.

It is important to note that significant care must be taken when using images downloaded from photo-sharing websites such as *Flickr*. As described in Section II-E, we went to great

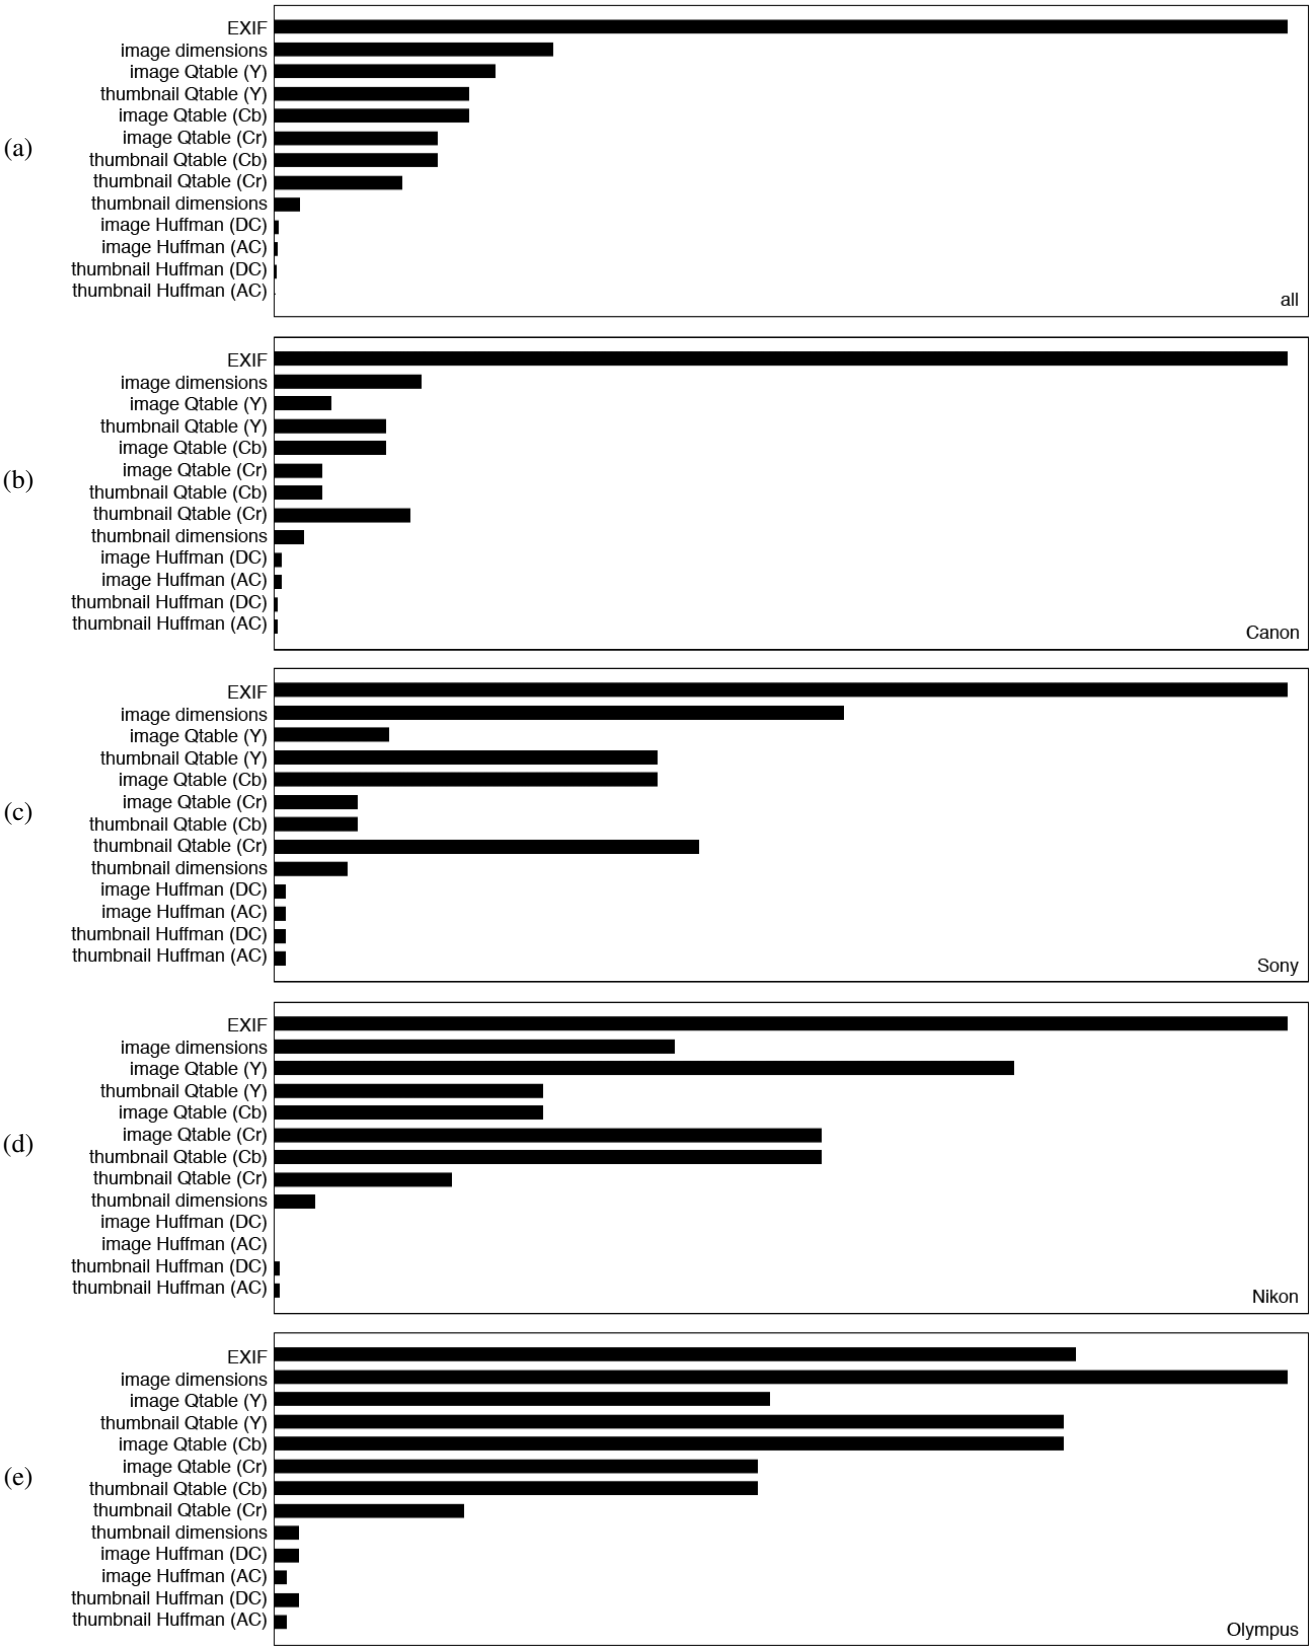

Fig. 5. Shown in panel (a) are the relative number of unique signature elements for all 9,163 camera configurations. The EXIF count is the most distinct and the Huffman codes are the least distinct. Shown below are the same statistics for the (b) Canon, (c) Sony, (d) Nikon, and (e) Olympus cameras. The Canon cameras are considerably more consistent than the other manufacturers. Each panel is normalized by the total number of camera configurations in the respective categories.

n	c	make	model
7	1	Canon	PS A540 | PS SD600 | PS SD630 | PS A530 | PS SD700 IS | PS SD550 | PS A700
7	2	Canon	PS SD1100 IS | PS S5 IS | PS A590 IS | PS SD870 IS | PS SX100 IS | PS SD850 IS | PS A720 IS
7	3	Canon	PS SD450 | PS A610 | PS SD400 | PS SD30 | PS SD430 Wireless | PS A620 | PS S80
7	4	Canon	PS SD450 | PS A610 | PS SD400 | PS SD550 | PS A620 | PS SD30 | PS S80
7	5	Canon	PS SX10 IS | PS S90 | PS G11 | PS A495 | PS SX1 IS | PS A3000 IS | PS A490
7	6	RIM	BB 8330 | BB 8310 | BB 8320 | BB 8130 | BB 8300 | BB 8330m | BB 8120
7	7	Sony	DSC-H50 | DSC-T300 | DSC-W150 | DSC-W220 | DSC-W170 | DSC-W230 | DSC-W300
7	8	Sony	DSC-H7 | DSC-T20 | DSC-W80 | DSC-H3 | DSC-H9 | DSC-T2 | DSC-T200
7	9	Sony	DSC-W55 | DSC-H5 | DSC-W35 | DSC-W70 | DSC-T30 | DSC-T10 | DSC-T50
7	10	Sony	DSC-W55 | DSC-H5 | DSC-W70 | DSC-T10 | DSC-W35 | DSC-T30 | DSC-T50
7	11	Sony	DSC-W55 | DSC-H5 | DSC-W70 | DSC-W35 | DSC-T10 | DSC-T30 | DSC-T50
8	1	Sony	DSC-H2 | DSC-W30 | DSC-W50 | DSC-T10 | DSC-W55 | DSC-W35 | DSC-W100 | DSC-H5
8	2	Sony	DSC-H9 | DSC-W90 | DSC-T100 | DSC-H7 | DSC-W200 | DSC-H3 | DSC-H10 | DSC-T20
8	3	Sony	DSC-T100 | DSC-W90 | DSC-H9 | DSC-H3 | DSC-H7 | DSC-T200 | DSC-W200 | DSC-H10
8	4	Sony	DSC-T700 | DSC-W80 | DSC-T20 | DSC-H9 | DSC-H7 | DSC-H3 | DSC-T2 | DSC-T100
8	5	Sony	DSC-W55 | DSC-T10 | DSC-H2 | DSC-W30 | DSC-W50 | DSC-H5 | DSC-W35 | DSC-T50
8	6	Sony	DSC-W80 | DSC-T100 | DSC-H3 | DSC-T2 | DSC-T700 | DSC-T20 | DSC-W90 | DSC-H9
8	7	Sony	DSC-W80 | DSC-T100 | DSC-H9 | DSC-T2 | DSC-T700 | DSC-H10 | DSC-T20 | DSC-H7
9	1	Canon	PS SD960 IS | PS SX20 IS | PS SD1300 IS | PS SD780 IS | PS D10 | PS SD940 IS | PS A1100 IS | PS SX200 IS | PS A2100 IS
9	2	Sony	DSC-H10 | DSC-T100 | DSC-T200 | DSC-H9 | DSC-W90 | DSC-H7 | DSC-H3 | DSC-W200 | DSC-T20
9	3	Sony	DSC-H9 | DSC-W90 | DSC-H7 | DSC-H3 | DSC-T100 | DSC-T2 | DSC-W200 | DSC-H10 | DSC-T20
10	1	Sony	DSC-H1 | DSC-T7 | DSC-V3 | DSC-W5 | DSC-P200 | DSC-T5 | DSC-P150 | DSC-T33 | DSC-T3 | DSC-M2
10	2	Sony	DSC-H1 | DSC-W5 | DSC-T7 | DSC-P200 | DSC-V3 | DSC-T33 | DSC-T5 | DSC-W150 | DSC-T3 | DSC-W7
10	3	Sony	DSC-H3 | DSC-H9 | DSC-T100 | DSC-W90 | DSC-H10 | DSC-W200 | DSC-T200 | DSC-H7 | DSC-T2 | DSC-T20
10	4	Sony	DSC-H9 | DSC-H7 | DSC-H10 | DSC-T100 | DSC-H3 | DSC-T200 | DSC-W90 | DSC-W200 | DSC-T20 | DSC-T2
11	1	Sony	DSC-W80 | DSC-T20 | DSC-T700 | DSC-H10 | DSC-T2 | DSC-H9 | DSC-H3 | DSC-T100 | DSC-W90 | DSC-H7 | DSC-W200
14	1	Sony	DSC-H1 | DSC-T7 | DSC-T1 | DSC-W1 | DSC-P100 | DSC-P200 | DSC-T33 | DSC-V3 | DSC-M2 | DSC-W5 | DSC-T3 | DSC-T5 | DSC-P150 | DSC-P93

Fig. 6. Each entry corresponds to an equivalence class of identical image, thumbnail and EXIF parameters. The first column is the size of the equivalence class n, the second column c is the number of equivalence classes of size n, the third and fourth columns are the camera make and models. The Canon Powershot model is abbreviated as PS, and the Research In Motion BlackBerry is abbreviated RIM BB.

lengths to ensure that the images being analyzed were not modified by a user or by *Flickr*. Even with these precautions, it is possible that some edited or altered images slipped through our system of checks. Anecdotally, we have noticed that JPEG images that are initially captured in RAW format and then converted, in software, to JPEG can be difficult to distinguish from images that are captured directly in JPEG format. Fortunately, it appears that many of these images can be filtered because RAW to JPEG converters often create images with resolutions that are different than the native camera resolutions. A more robust system would require determining which JPEG quantization tables and Huffman codes are used by popular RAW conversion software, and then eliminating any images that employ them.

We have focused on the JPEG standard because it remains the most ubiquitous image format employed by digital cameras. We note that should JPEG 2000 become more widely used, our basic approach could be extended to this format. As with the JPEG standard, the JPEG 2000 standard contains the same basic parameters that we have analyzed: an EXIF header, thumbnail, Huffman coding, and quantization (of the wavelet, as opposed to DCT, coefficients).

The power of our forensic analysis lies in the ability to acquire signatures from a wide variety of cameras and cellphones. This poses significant challenges as new cameras and cellphones are constantly released. We expect to continue building our database of images and camera information in order to keep up with these continual changes.

ACKNOWLEDGMENT

This work was supported by a gift from Adobe Systems, Inc. and a gift from Microsoft, Inc.

REFERENCES

[1] J. Fridrich, D. Soukal, and J. Lukáš, "Detection of copy move forgery in digital images," in *Proceedings of Digital Forensic Research Workshop*, August 2003.

[2] A. Popescu and H. Farid, "Exposing digital forgeries by detecting duplicated image regions," Department of Computer Science, Dartmouth College, Tech. Rep. TR2004-515, 2004.

[3] S. Bayram, H. T. Sencar, and N. Memon, "A survey of copy-move forgery detection techniques," in *IEEE Western New York Image Processing Workshop*, 2008.

[4] X. Pan and S. Lyu, "Region duplication detection using image feature matching," *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 4, pp. 857 –867, 2010.

[5] A. Popescu and H. Farid, "Exposing digital forgeries by detecting traces of re-sampling," *IEEE Transactions on Signal Processing*, vol. 53, no. 2, pp. 758–767, 2005.

[6] M. Kirchner and T. Gloe, "On resampling detection in re-compressed images," in *IEEE Workshop on Information Forensics and Security*, 2009, pp. 21–25.

[7] A. C. Popescu and H. Farid, "Exposing digital forgeries in color filter array interpolated images," *IEEE Transactions on Signal Processing*, vol. 53, no. 10, pp. 3948–3959, 2005.

[8] M. Kirchner, "Efficient estimation of CFA pattern configuration in digital camera images," in *SPIE Conference on Media Forensics and Security*, 2010.

[9] Z. Lin, R. Wang, X. Tang, and H.-V. Shum, "Detecting doctored images using camera response normality and consistency," in *Computer Vision and Pattern Recognition*, San Diego, CA, 2005.

[10] M. Johnson and H. Farid, "Exposing digital forgeries in complex lighting environments," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 2, pp. 450–461, 2007.

[11] W. Zhang, X. Cao, J. Zhang, J. Zhu, and P. Wang, "Detecting photographic composites using shadows," in *IEEE International Conference on Multimedia and Expo*, 2009, pp. 1042–1045.

[12] C. Riess and E. Angelopoulou, "Scene illumination as an indicator of image manipulation," in *12th International Workshop on Information Hiding*, 2010.

[13] E. Kee and H. Farid, "Exposing digital forgeries from 3-D lighting environments," in *Workshop on Information Forensics and Security*, 2010.

[14] M. Johnson and H. Farid, "Exposing digital forgeries through chromatic aberration," in *ACM Multimedia and Security Workshop*, Geneva, Switzerland, 2006.

[15] T. Gloe, A. Winkler, and K. Borowka, "Efficient estimation and large-scale evaluation of lateral chromatic aberration for digital image forensics," in *SPIE Conference on Media Forensics and Security*, 2010.

[16] J. Lukáš, J. Fridrich, and M. Goljan, "Digital camera identification from sensor noise," *IEEE Transactions on Information Security and Forensics*, vol. 1, no. 2, pp. 205–214, 2006.

[17] M. Chen, J. Fridrich, M. Goljan, and J. Lukas, "Determining image origin and integrity using sensor noise," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 1, pp. 74–90, 2008.

[18] M. Stamm and K. J. R. Liu, "Forensic detection of image manipulation using statistical intrinsic fingerprints," *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 3, pp. 492–506, 2010.

[19] H. Farid, "A survey of image forgery detection," *IEEE Signal Processing Magazine*, vol. 2, no. 26, pp. 16–25, 2009.

[20] A. Popescu and H. Farid, "Statistical tools for digital forensics," in *6th International Workshop on Information Hiding*, Toronto, Canada, 2004.

[21] J. Zhang, H. Wang, and Y. Su, "Detection of double-compression in JPEG2000 images," in *International Symposium on Intelligent Information Technology Application*, vol. 1, 2008, pp. 418–421.

[22] Z. Qu, W. Luo, and J. Huang, "A convolutive mixing model for shifted double JPEG compression with application to passive image authentication," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2008, pp. 1661–1664.

[23] M. Sorell, "Conditions for effective detection and identification of primary quantization of re-quantized JPEG images," in *International Conference on Forensic Applications and Techniques in Telecommunications, Information, and Multimedia*, 2008.

[24] B. Mahdian and S. Saic, "Detecting double compressed JPEG images," in *3rd International Conference on Imaging for Crime Detection and Prevention*, 2009.

[25] P. Alvarez, "Using extended file information (EXIF) file headers in digital evidence analysis," *International Journal of Digital Evidence*, vol. 2, no. 3, 2004.

[26] H. Farid, "Digital image ballistics from JPEG quantization," Department of Computer Science, Dartmouth College, Tech. Rep. TR2006-583, 2006.

[27] J. Kornblum, "Using JPEG quantization tables to identify imagery processed by software," *Digital Investigation*, vol. 5, pp. S21–S25, 2008.

[28] H. Farid, "Digital image ballistics from JPEG quantization: A followup study," Department of Computer Science, Dartmouth College, Tech. Rep. TR2008-638, 2008.

[29] E. Kee and H. Farid, "Digital image authentication from thumbnails," in *SPIE Symposium on Electronic Imaging*, San Jose, CA, 2010.

[30] "Digital compression and coding of continuous-tone still images, Part 1: Requirements and guidelines," ISO/IEC JTC1 Draft International Standard 10918-1, 1991.

[31] G. Wallace, "The JPEG still picture compression standard," *IEEE Transactions on Consumer Electronics*, vol. 34, no. 4, pp. 30–44, 1991.

[32] "Exchangeable image file format for digital still cameras: Exif version 2.2," Japan Electronics and Information Technology Industries Association, 2002.

[33] M. Goljan, J. Fridrich, and T. Filler, "Large scale test of sensor fingerprint camera identification," in *SPIE Conference on Media Forensics and Security*, 2009.

[34] ——, "Managing a large database of camera fingerprints," in *SPIE Conference on Media Forensics and Security*, 2010.

[35] S. Bayram, H. T. Sencar, and N. Memon, "Efficient techniques for sensor fingerprint matching in large image and video databases," in *SPIE Conference on Media Forensics and Security*, 2010.

APPENDIX I

Two quantization tables employed by *Flickr*. Shown from left to right are 8×8 quantization tables corresponding to the luminance (Y), and chrominance (Cb and Cr) channels. The values correspond to the quantization applied to DCT coefficients, Equation (2).

6	5	6	6	7	10	20	29	7	7	10	19	40	40	40	40
4	5	5	7	9	14	26	37	7	8	10	26	40	40	40	40
4	6	6	9	15	22	31	38	10	10	22	40	40	40	40	40
6	8	10	12	22	26	35	39	19	26	40	40	40	40	40	40
10	10	16	20	27	32	41	45	40	40	40	40	40	40	40	40
16	23	23	35	44	42	48	40	40	40	40	40	40	40	40	40
20	24	28	32	41	45	48	41	40	40	40	40	40	40	40	40
24	22	22	25	31	37	40	40	40	40	40	40	40	40	40	40

1	1	1	1	1	2	4	6	1	1	2	4	8	8	8	8
1	1	1	1	2	3	5	7	1	2	2	5	8	8	8	8
1	1	1	2	3	4	6	8	2	2	4	8	8	8	8	8
1	2	2	2	4	5	7	8	4	5	8	8	8	8	8	8
2	2	3	4	5	6	8	9	8	8	8	8	8	8	8	8
3	5	5	7	9	8	10	8	8	8	8	8	8	8	8	8
4	5	6	6	8	9	10	8	8	8	8	8	8	8	8	8
5	4	4	5	6	7	8	8	8	8	8	8	8	8	8	8