

Compression Forensics Beyond The First Significant Digit

Sujoy Chakraborty and Matthias Kirchner

Department of Electrical and Computer Engineering, Binghamton University, Binghamton, NY 13902, USA

Abstract

We study characteristics of the second significant digits of block-DCT coefficients computed from digital images. Following previous works on compression forensics based on first significant digits, we examine the merits of stepping towards significant digits beyond the first. Our empirical findings indicate that certain block-DCT modes follow Benford's law of second significant digits extremely well, which allows us to distinguish between never-compressed images and decompressed JPEG images even for the highest JPEG compression quality of 100. As for multiple-compression forensics we report that second significant digit histograms are highly informative on their own, yet cannot further improve already good performances of classification schemes that work with first significant digits alone.

Introduction

The wide-spread use of the JPEG compression standard for still images makes the analysis of JPEG images a major branch of media forensics [1]. Most digital cameras store images in JPEG format, typically using customized compression or file format settings [2, 3], and most likely images will be stored as JPEG again using a different set of compression settings after processing or manipulation of any sort. This introduces various forms of requantization artifacts, which may be exploited to infer the compression history of a given image. Scenarios vary greatly, including for instance the detection and characterization of previous JPEG compressions of images stored in bitmap formats [4, 5], the detection and characterization of multiple JPEG compressions [6–11], possibly in the presence of other forms of processing in between consecutive compression steps [12, 13], or the detection of local image manipulations [9, 14–16].

Among the most successful approaches for multiple compression forensics are those that rely on statistics of the first significant digits (FSDs) of discrete cosine transform (DCT) coefficients. It has been observed that block-DCT coefficients of natural images follow (a generalized) Benford's law of first significant digits [17, 18], while lossy JPEG compression changes the FSD distribution of DCT coefficients. The analysis of empirical FSD histograms is effectively equivalent to aggregating statistics over DCT coefficient histograms. Hence, FSD-based approaches promise more compact feature representations, yet they cannot be expected to fundamentally outperform forensic techniques that work with first-order DCT coefficient statistics directly [19]. Indeed, the literature has demonstrated that FSD-based classifiers are highly capable of identifying the number of compression cycles an image underwent in a relatively low-dimensional feature space [10].

It has been noted only recently that also the second significant digits (SSDs) of block-DCT coefficients exhibit a highly regular behavior across different image databases [20], partially in correspondence with Benford's law. Here, we follow this path and

explore to what degree peculiarities in the distribution of block-DCT SSDs can be exploited for compression forensics. While it cannot be expected that switching from first significant digits to higher-order digits will result in tremendous performance gains, it is our hope to deepen the understanding of how significant digits of DCT coefficients behave in various forensically relevant settings, and whether there exist scenarios for which it is beneficial to consider also significant digits beyond the first. Before we delve into our exploratory analysis, the following two sections briefly summarize how Benford's law characterizes the empirical distribution of significant digits and how FSD features have been utilized in prior work on compression forensics.

Significant Digits and Benford's Law

Denoting $s(x)$ as the decimal¹ significand of a non-zero real number $x \in \mathbb{R} \setminus \{0\}$, $s(x) = 10^{\log|x| - \lfloor \log|x| \rfloor}$, the s -th significant digit of x , $d_s(x)$, $s \in \mathbb{N}$, is given as

$$d_s(x) = \left\lfloor 10^{s-1} s(x) \right\rfloor - 10 \left\lfloor 10^{s-2} s(x) \right\rfloor. \quad (1)$$

By definition, the first significant digit (FSD) of $x \neq 0$ is never zero, $d_1(x) \in [1..9]$. Significant digits $d_s(x)$ with $s > 1$ can also take on value zero. Benford's law [21, 22] concerns the statistical distribution of significant digits. It has been found to apply to various types of synthetic and empirical data. Denoting $\Pr(D_s = d_s)$ as the probability that the s -th significant digit equals d_s , Benford's law states that [22]

$$\Pr((D_1, D_2, \dots, D_r) = (d_1, d_2, \dots, d_r)) = \log \left(1 + \left(\sum_{s=1}^r 10^{r-s} d_s \right)^{-1} \right). \quad (2)$$

Specifically, this implies for first and second significant digits that

$$\Pr(D_1 = d_1) = \log(1 + 1/d_1) \quad \text{and} \quad (3)$$

$$\Pr(D_2 = d_2) = \sum_{d_1=1}^9 \log(1 + 1/(10d_1 + d_2)), \quad (4)$$

respectively. Figure 1 gives a graphical representation of the two probability mass functions. Observe that the SSD distribution is much more uniform than the FSD distribution. Generally speaking, it follows from Benford's law that the marginal distribution of the s -th significant digit approaches the uniform distribution as $s \rightarrow \infty$ [23]. A sufficient condition for Benford's law to be satisfied is a uniform distribution of $\log s(x)$ over the interval $[0, 1)$ [22].

FSD-Based JPEG Compression Forensics

First significant digits of block-DCT modes of natural images are commonly assumed to obey (a generalized variant of) Benford's law [17, 18]. A common working assumption of forensic

¹We work with logarithms to base 10 in this paper.




Figure 1. Benford's law for first (FSD) and second (SSD) significant digits.

techniques is that FSD distributions after lossy JPEG compression exhibit a fundamentally different behavior. By analyzing empirical FSD histograms, this can be exploited to detect traces of prior JPEG compressions in decompressed images or to determine the number and parameters of previous JPEG compression cycles. Early methods leaned towards explicit tests for Benford's law to be satisfied [18, 24]. More recent techniques draw more heavily on machine learning support. Most recently, Milani et al. [10] determined empirically a set of FSD features for multiple compression forensics by assessing a large number of possible digit combinations from a predetermined set of nine low-frequency JPEG coefficients proposed in [25]. Specifically, the authors focussed on fixed n -tuples of digits and considered the corresponding $n \times 9$ FSD histogram bins as feature space. Digits 2, 5 and 6 were found to work particularly well. The classification algorithm combines a set of binary classifiers, individually trained to distinguish between various compression settings one at a time. The final classification result is then aggregated from the combined binary decisions.

Experiments With Block-DCT SSDs

Positive reports about FSD-based compression forensics in the literature have also led to a number of counter-forensic techniques [27] that attempt to restore block-DCT coefficient FSD histograms after JPEG compression [28, 29]. Yet we demonstrated in a recent work [20] that restoration algorithms under a minimum cost constraint can lead to detectable artifacts in the histograms of second significant digits. One of the major findings along the way was that block-DCT SSD histograms from never-compressed images exhibit a highly regular behavior. A question that naturally arises is thus whether SSD histogram features are similarly suitable for compression forensics as FSD features.

Experimental Setup

We work with the 1338 UCID images [26] in our experiments. All images are of size 384×512 . The images were converted to grayscale before further processing.² We use the Independent JPEG Group reference library with floating-point DCT implementation and standard quantization tables to obtain JPEG versions of the database, considering single compression with quality factors from the set $\mathcal{Q} = \{35, 40, \dots, 100\}$ and double compression with quality factor combinations in $\mathcal{Q} \times \mathcal{Q}$. For a given tuple $\mathbf{q} = (q_1, q_2) \in \mathcal{Q} \times \mathcal{Q}$, we use notation $q_1 \rightarrow q_2$ to refer to compression with quality factor q_1 followed by a second compression

²ImageMagick convert with option `-grayscale Rec601Luma`.

with quality factor q_2 . Similar to Milani et al. [10], we work with a simplified setup that limits the set of primary quality factors to

$$\mathcal{Q}_{q_2}^\Delta = \{q_1 \in \mathcal{Q} : 0 < |q_2 - q_1| \leq \Delta\}, \quad (5)$$

i. e., for a specific choice of q_2 , we consider only quality factor combinations $(q_1, q_2) \in \mathcal{Q}_{q_2}^\Delta \times \{q_2\}$.

We analyze DCT coefficients, rounded to six digits, which are computed from non-overlapping 8×8 pixel blocks $(y_{m,n})$, $0 \leq m, n \leq 7$, with integer intensities $y_{m,n} \in [0..255]$ as

$$x_{i,j} = \sum_{m=0}^7 \sum_{n=0}^7 (y_{m,n} - 128) \cdot b_{m,n}^{(i,j)}. \quad (6)$$

The elements of the (i, j) -th basis vector $\mathbf{b}^{(i,j)}$ are given as

$$b_{m,n}^{(i,j)} = \frac{c_i c_j}{4} \cdot \cos\left(\frac{\pi i (2m+1)}{16}\right) \cdot \cos\left(\frac{\pi j (2n+1)}{16}\right), \quad (7)$$

where $0 \leq i, j \leq 7$, $c_0 = 1/\sqrt{2}$ and $c_i = 1$ for $i > 0$. Pixel blocks are aligned with the JPEG grid, if the image was previously stored as JPEG. Normalized FSD and SSD histograms are computed per DCT mode (i, j) . Omitting index (i, j) for the sake of brevity, and denoting $\mathbf{x} = (x_k)$, $0 \leq k < N$, as the vector of coefficients satisfying $\{x : x \neq 0 \wedge x \neq d_1(x)\}$, the s -th digit normalized histogram is

$$h_s(d, \mathbf{x}) = \frac{1}{N} \sum_{k=0}^{N-1} \delta(d - d_s(x_k)), \quad (8)$$

where $\delta(\cdot)$ denotes the Kronecker delta function.

Never-compressed Images

Figure 2 illustrates the distribution of second significant digits obtained from block-DCT modes of never-compressed images [20]. Each of the 8×8 sub-graphs depicts the median SSD histogram (aggregated over all images in the UCID database) in correspondence to the DCT coefficient index. Plots of the ‘‘ideal’’ distribution according to Benford's law in Equation (4) are given as reference with each empirical histogram. The figure indicates that the majority of DCT modes adhere to Benford's law strikingly well. There are three strong outliers at DCT coefficient indices $(i, j) \in \{(4, 0), (0, 4), (4, 4)\}$. This can be explained by the special form of the DCT for those frequencies. As pointed out in [30] in the context of steganography, it is straightforward to verify that Equation (7) evaluates to $b_{m,n}^{(i,j)} = \pm \frac{1}{8}$ for all m, n when $(i, j) \in \{(4, 0), (0, 4), (4, 4)\}$. This implies that DCT coefficients computed from integer pixel intensities will be exclusively integer multiples of $\frac{1}{8}$ at those frequencies, i. e., second significant digits $\{0, 4, 9\}$ can only occur for coefficients $\{x : |x| > 10\}$. Considering the Laplacian-like distribution of DCT coefficients, these SSD histogram bins will be populated only minimally. Similar but by far less strong artifacts can also be observed for other coefficient indices with even row and column indices [20].

Decompressed JPEG Images

Figure 3 continues with a closer look at block-DCT SSD distributions after JPEG compression and decompression. Specifically, we plot the empirical histograms aggregated over all UCID images after compression with quality factor 100. The histograms




Figure 2. Empirical SSD distributions of 8×8 block-DCT coefficients from the UCID image database. Subfigures are arranged in correspondence with DCT mode indices and depict the medians, the 25% quartiles and the 75% quartiles of relative frequencies for digits $\{0, \dots, 9\}$ (from left to right within each subfigure). Blue dots denote the distribution according to Benford's law (equal for all coefficients, see Figure 1). All subfigures are plotted on the same vertical scale.

are clearly less uniform than those from never-compressed images in Figure 2, and it is worth pointing out the consistent bias towards SSDs 0 and 9. This is to be expected, considering that JPEG compression with quality factor 100 corresponds to plain rounding of all DCT coefficients. Spatial domain rounding and truncation errors in the decompression step will yield non-integer DCT coefficients in Equation (6). The resulting coefficients are generally distributed closely around the corresponding dequantized (integer-valued) JPEG coefficients [4], with a strong bias towards the “ends” of the SSD histogram. Overall, already a visual comparison of Figures 2 and 3 thus suggests that block-DCT SSD histograms may be a valuable source of information for compression forensics.

For a quantitative analysis, we compute χ^2 -differences between empirical histograms and the SSD distribution according to Benford's law,

$$\chi^2 = \sum_{d_2=0}^9 \frac{\left(h_2(d_2, \mathbf{x}) - \sum_{d_1=1}^9 \log(1 + 1/(10d_1 + d_2)) \right)^2}{\sum_{d_1=1}^9 \log(1 + 1/(10d_1 + d_2))}. \quad (9)$$

We evaluate Equation (9) for each of the 16 DCT coefficients with odd row and column indices, $(i, j) = (2w + 1, 2v + 1)$, $0 \leq v, w \leq 3$, and take the average χ^2 as aggregated measure per image. Figure 4 illustrates the results of this procedure by reporting the minimum and median average χ^2 over all images in the database after JPEG compression with quality factors $q_1 \in \mathcal{Q}$. The horizontal line corresponds to the maximum average χ^2 over all never-compressed images in the database, indicating that it is possible to perfectly




Figure 3. Empirical SSD distributions of 8×8 block-DCT coefficients from the UCID image database after JPEG compression with quality factor 100. Subfigures are arranged in correspondence with DCT mode indices and depict the medians, the 25% quartiles and the 75% quartiles of relative frequencies for digits $\{0, \dots, 9\}$ (from left to right within each subfigure). Blue dots denote the distribution according to Benford's law (equal for all coefficients, see Figure 1). All subfigures are plotted on the same vertical scale.

separate between never-compressed images and previously JPEG-compressed images across all tested JPEG quality factors. We emphasize that FSD histograms are generally not as discriminative for images compressed with extremely high quality factors. An FSD-based JPEG compression detector adopting above strategy with zero false negatives would result in a 75% missed detection rate for $q_1 = 100$, for instance.

Double Compression

Differences between never-compressed and compressed images become generally more evident after stronger compression. A more interesting question is to what degree SSD histograms from single-compressed and double-compressed JPEG images differ. Figure 5 sheds some light on this aspect by comparing SSD histograms from images compressed once with JPEG quality factor 70, and images compressed twice with a sequence of quality factors $80 \rightarrow 70$. Overall, observe that the strong quantization of lower JPEG quality factors will tend to produce relatively uniform SSD histograms, especially so for higher frequencies. Differences between the two exemplary settings are still clearly noticeable in certain DCT modes, yet it can be expected that a clear distinction may not always be possible.

Classifiers and Feature Selection We train binary soft-margin SVM classifiers with RBF kernel to distinguish between single-compressed images (compressed once with quality factor q_2) and double-compressed images (compressed twice with $q_1 \rightarrow q_2$, $q_1 \in$

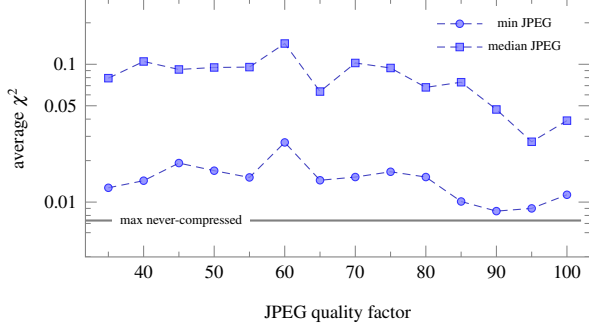


Figure 4. χ^2 -differences between empirical SSD distributions and Benford’s law from decompressed JPEG images (average χ^2 , aggregated over odd-indexed DCT coefficients). The graphs show the minimum and median differences across all JPEG-compressed UCID images. The horizontal line denotes the maximum averaged χ^2 in the set of never-compressed images.

$\mathcal{Q}_{q_2}^\Delta$). One classifier is trained per secondary JPEG quality factor q_2 (which is available from the query image). Training data for both classes are compiled from the same images with a 50/50 training/validation split of the database. Both training sets are of equal size; the double-compressed set comprises a randomized mixture of quality factors $\mathcal{Q}_{q_2}^\Delta \times \{q_2\}$. The trained classifier is then fed the remaining half of images in the validation step, each image compressed once with q_2 , and double-compressed in total $|\mathcal{Q}_{q_2}^\Delta|$ times with all combinations in $\mathcal{Q}_{q_2}^\Delta \times \{q_2\}$. We report average classification accuracies as performance measure, aggregated over ten randomized hold-out training/validation splits of the database.

Feature vectors are constructed from empirical normalized histogram bins (d, i, j) , corresponding to significant digit d from the (i, j) -th DCT mode. It will be evident from the context whether we refer to FSD or SSD features. We generally do not consider digit $d = 9$, since normalized histogram bins always add up to unity, implying that one of the bins is implicitly given by all others. Our selection of features is guided by the Fisher score [31], computed individually for each feature index (d, i, j) from all images in the database. Specifically, we consider for each combination of quality factors $q_1 \rightarrow q_2$, $q_1 \in \mathcal{Q}_{q_2}^\Delta$, two sets of images, \mathcal{X}_{q_2} and $\mathcal{X}_{q_1 \rightarrow q_2}$. The indices of those sets are indicative of the images’ compression history. Denote $\mu_k^{(d,i,j)}$ and $\sigma_k^{(d,i,j)}$ the empirical mean and standard deviation over all histogram bins (d, i, j) in image set \mathcal{X}_k , $k \in \{q_2, q_1 \rightarrow q_2\}$, the Fisher score for a specific feature index is

$$F_{q_1, q_2}(d, i, j) = \frac{\sum_k (\mu_k^{(d,i,j)} - \mu^{(d,i,j)})^2}{\sum_k (\sigma_k^{(d,i,j)})^2}, \quad (10)$$

with $\mu^{(d,i,j)}$ denoting the empirical mean over both classes. Higher scores indicate better separability between single-compressed and double-compressed images based on feature (d, i, j) . An N -dimensional feature set is obtained per quality factor q_2 by selecting the N features that appear most frequently in the top-100 Fisher score ranking of all quality factor combinations $(q_1, q_2) \in \mathcal{Q}_{q_2}^\Delta \times \{q_2\}$.

Classification Results Figure 6 reports average classification accuracies per JPEG compression quality factor q_2 for the distinction between single compression with quality factor q_2 and




Figure 5. Empirical SSD distributions of 8×8 block-DCT coefficients from the UCID image database after JPEG compression with quality factor 70 and double compression with quality factors $80 \rightarrow 70$. Subfigures are arranged in correspondence with DCT mode indices and depict the medians of relative frequencies for digits $\{0, \dots, 9\}$ (from left to right within each subfigure). Blue dots denote the distribution according to Benford’s law (equal for all coefficients, see Figure 1). All subfigures are plotted on the same vertical scale.

double compression with quality factors $q_1 \rightarrow q_2$, $q_1 \in \mathcal{Q}_{q_2}^\Delta$, when $\Delta = 10$. We consider three different feature sets comprising SSD histogram bins, FSD histogram bins, or a combination of both, respectively. All feature sets are of dimension 60 and were found using the procedure described above. The results indicate that the classification accuracy is generally very high in this scenario, except for quality factor combinations where $\{100\} \cap \mathcal{Q}_{q_2}^\Delta \neq \emptyset$. A comparison with Figure 7, which parallels the setting above, but excludes double compressions with $q_1 = 100$, confirms that $q_1 = 100$ is indeed a weak spot of all feature sets. This can be expected, since the difference between never-compressed images and images compressed with highest quality is marginal. It is advised to use more specialized techniques for such scenarios, especially also when multiple compressions with the same quality factor are concerned [11]. More interestingly in our context, observe that SSD features are inferior to FSD features for lower quality factors. This can be explained from the combination of a Laplacian-like distribution of DCT coefficients and increasingly stronger quantization factors, which implies that SSDs of DCT coefficients from decompressed images in Equation (6) will be dominated more strongly by spatial domain rounding and truncation noise as the compression quality decreases.

A closer look at the selected features indicates that the analysis of double compression with a relatively high secondary quality factor q_2 can leverage digits from a wider range of DCT coefficients than it is the case with lower quality factors. Figures 8 and 9 illustrate this effect by visualizing the selected SSD and FSD histogram bins for $q_2 = 90$ and $q_2 = 70$ in the DCT coef-




Figure 6. Average classification accuracies for single ($q_2 \in \mathcal{Q}$) vs. double compression ($q_1 \rightarrow q_2, q_1 \in \mathcal{Q}_{q_2}^\Delta$) detection, $\Delta = 10$. Feature space dimension 60.




Figure 7. Average classification accuracies for single ($q_2 \in \mathcal{Q}$) vs. double compression ($q_1 \rightarrow q_2, q_1 \in \mathcal{Q}_{q_2}^\Delta \setminus \{100\}$) detection, $\Delta = 10$. Feature space dimension 60.

ficient plane. It is worth pointing out that the SSD features for $q_2 = 70$ are in good alignment with the differences that can be observed in the distribution plots in Figure 5. We also note that that SSD features for higher quality factors tend to include more low-frequency DCT coefficients than we see with FSD features. For a better understanding of the different classification accuracies for lower quality factors, Figure 10 compares the FSD, SSD and combined FSD/SSD feature selections for quality factor $q_2 = 60$ in a single visualization. The graph indicates that both FSD and SSD features draw exclusively from low-frequency coefficients. The combined feature set, which yields an accuracy comparable to features compiled from FSDs alone, is even more “strict” in this regard, at the same time it favors mostly FSD features (red circles) over SSD features (blue squares).

As far as the dimension of the feature space is concerned, we observed only a relatively small influence on the average classification accuracy in the setup described above. Figure 11 exemplarily reports average accuracies obtained with different numbers of features for $q_2 = 70$. More importantly, it is has to be pointed out that the classification accuracy drops considerably under more realistic scenarios when larger differences between primary and secondary compression qualities have to be considered. Figure 12 gives an example for $\Delta = 20$. The training set in such more diverse settings contains less instances of specific quality factor combinations, which makes binary classification generally harder. A similar effect can be expected for a higher (and more practically relevant) granularity of quality factors under consideration. Training sets of

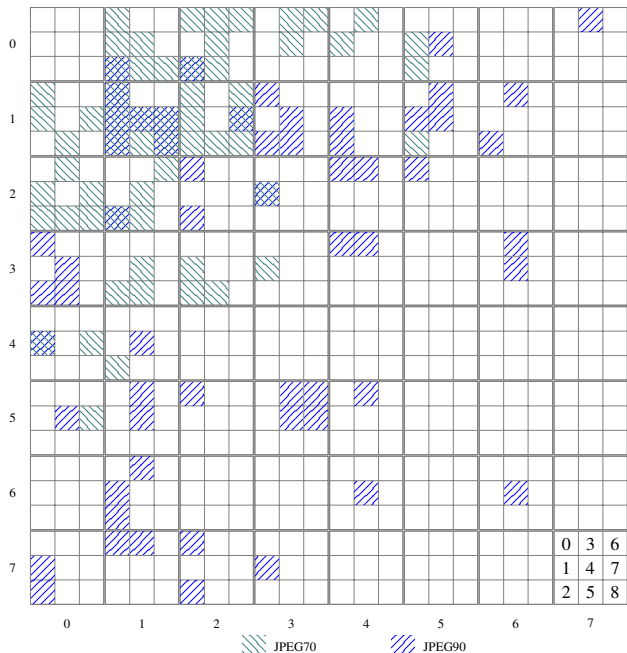


Figure 8. SSD features for double compression detection with $q_2 = 70$ and $q_2 = 90$, $\Delta = 10$. Each DCT mode is represented by a 3×3 sub-block, index (0, 0) in the upper left corner is the DC mode; sub-blocks encode digits $d \in [0..8]$ (arranged in column major order). Highlighted digits contribute to the feature space.

significantly increased size and/or multi-class classifiers are more advisable in such settings [10, 25].

Concluding Remarks

We have explored characteristics of the second significant digits (SSDs) of block-DCT coefficients in the context of JPEG compression forensics. While distributions of first significant digits are understood relatively well [17, 18], little had been known about significant digits beyond the first. We have focused here on relatively simple scenarios to put most emphasis on a better understanding of the differences between first and second significant digits. Our empirical findings indicate that never-compressed images adhere to Benford’s law extremely well as far as SSDs of odd-indexed DCT coefficients are concerned. SSDs computed from decompressed images are highly sensitive to prior compression, which allowed us to perfectly distinguish even between never-compressed images and images that underwent JPEG compression with highest quality (quality factor 100). This is not possible based on first significant digit (FSD) characteristics. As for the distinction between single and double compression, we observed that SSD-based features are by and large comparable to features obtained from FSD histograms. Both can yield very high classification accuracies in simple test scenarios, yet FSD features tend to be more robust against lower compression qualities. Overall, we find that our results are in good alignment with the results reported by Milani et al. [10] based on *quantized* JPEG coefficients. Along these lines, we mention that the analysis of SSDs from quantized DCT coefficients is largely unfeasible due to the limited support of coefficient histograms even already after moderate compression. In summary,




Figure 9. FSD features for double compression detection with $q_2 = 70$ and $q_2 = 90$, $\Delta = 10$. Each DCT mode is represented by a 3×3 sub-block, index (0,0) in the upper left corner is the DC mode; sub-blocks encode digits $d \in [0..8]$ (arranged in column major order). Highlighted digits contribute to the feature space. (Digit $d = 0$ is only included for consistency with SSD-feature plots.)

it thus seems questionable that there are profound reasons to step beyond the first significant digit for basic compression forensics purposes. Nevertheless, it should be understood that higher-order significant digits *are* highly informative, especially also in the presence of counter-forensics [20]. In a broader context, we surmise that the observation and characterization of “non-Benford” SSD distribution peculiarities due to the discreteness of digital images is an excellent stepping stone towards more accurate models of certain DCT modes. We leave this to future work, along with an exploration of possible applications to image forensics.

References

- [1] R. Böhme and M. Kirchner, “Media forensics,” in *Information Hiding*, S. Katzenbeisser and F. Petitcolas, Eds. Artech House, 2016, ch. 9, pp. 231–259.
- [2] E. Kee, M. K. Johnson, and H. Farid, “Digital image authentication from JPEG headers,” *IEEE Transactions on Information Forensics and Security*, vol. 6, no. 3, pp. 1066–1075, 2011.
- [3] T. Gloe, “Forensic analysis of ordered data structures on the example of JPEG files,” in *IEEE Workshop on Information Forensics and Security (WIFS)*, 2012, pp. 139–144.
- [4] Z. Fan and R. L. de Queiroz, “Identification of bitmap compression history: JPEG detection and quantizer estimation,” *IEEE Transactions on Image Processing*, vol. 12, no. 2, pp. 230–235, 2003.
- [5] R. Neelamani, R. de Queiroz, Z. Fan, S. Dash, and R. G. Baraniuk, “JPEG compression history estimation for color images,” *IEEE Transactions on Image Processing*, vol. 15, no. 6, pp. 1365–1378, 2006.
- [6] J. Lukáš and J. Fridrich, “Estimation of primary quantization matrix

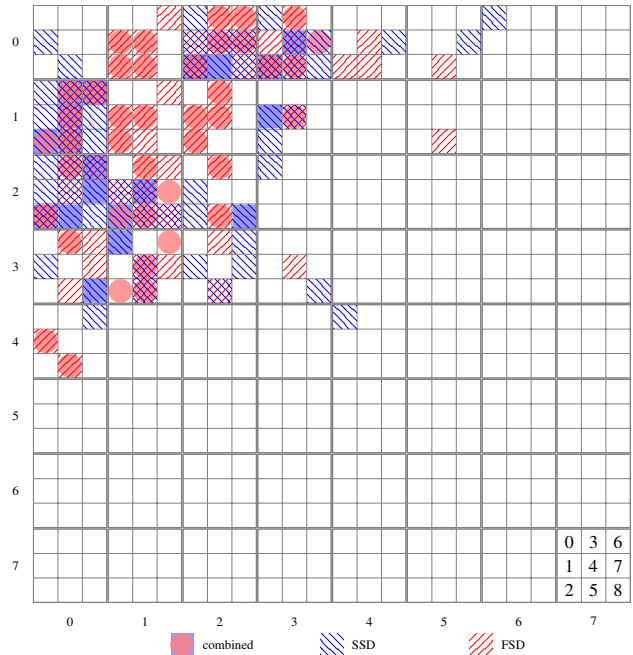


Figure 10. Combined FSD/SSD features for double compression detection with $q_2 = 60$, $\Delta = 10$. Each DCT mode is represented by a 3×3 sub-block, index (0,0) in the upper left corner is the DC mode; sub-blocks encode digits $d \in [0..8]$ (arranged in column major order). Highlighted digits contribute to the feature space. Red denotes FSDs, blue stands for SSDs. DCT coefficients can contribute both FSD and SSD features.

in double compressed JPEG images,” in *Digital Forensic Research Workshop*, 2003.

- [7] A. C. Popescu and H. Farid, “Statistical tools for digital forensics,” in *Information Hiding, 6th International Workshop*, ser. Lecture Notes in Computer Science, J. Fridrich, Ed., vol. 3200, 2004, pp. 128–147.
- [8] F. Huang, J. Huang, and Y. Q. Shi, “Detecting double JPEG compression with the same quantization matrix,” *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 4, pp. 848–856, 2010.
- [9] W. Luo, J. Huang, and G. Qiu, “JPEG error analysis and its applications to digital image forensics,” *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 3, pp. 480–491, 2010.
- [10] S. Milani, M. Tagliasacchi, and S. Tubaro, “Discriminating multiple JPEG compressions using first digit features,” *APSIPA Transactions on Signal and Information Processing*, vol. 3, e19, 2014.
- [11] M. Carnein, P. Schöttle, and R. Böhme, “Forensics of high-quality JPEG images with color subsampling,” in *IEEE International Workshop on Information Forensics and Security (WIFS)*, 2015.
- [12] P. Ferrara, T. Bianchi, A. D. Rosa, and A. Piva, “Reverse engineering of double compressed images in the presence of contrast enhancement,” in *IEEE International Workshop on Multimedia Signal Processing (MMSP)*, 2013, pp. 141–146.
- [13] V. Conotter, P. Comesaña, and F. Pérez-González, “Forensic detection of processing operator chains: Recovering the history of filtered JPEG images,” *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 11, pp. 2257–2269, 2015.
- [14] H. Farid, “Exposing digital forgeries from JPEG ghosts,” *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 1, pp. 154–160, 2009.
- [15] Z. Lin, J. He, X. Tang, and C.-K. Tang, “Fast, automatic and fine-

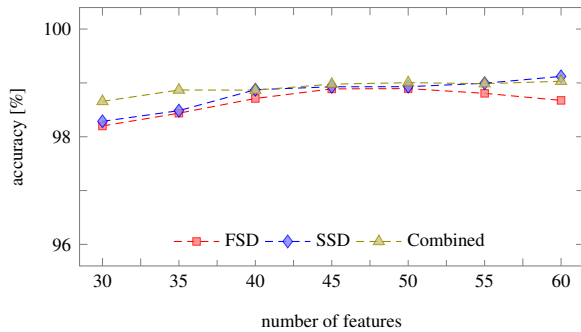


Figure 11. Average classification accuracy for single ($q_2 = 70$) vs. double compression ($q_1 \rightarrow 70, q_1 \in \mathcal{Q}_{70}^\Delta$) detection with different feature space dimensionalities, $\Delta = 10$.

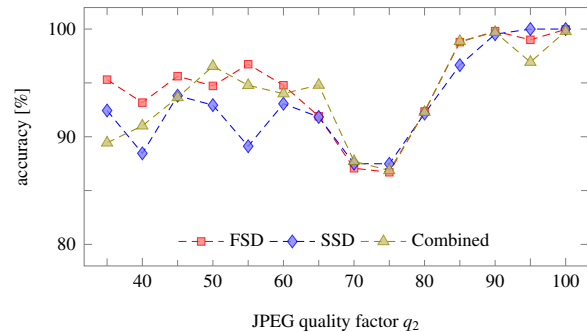


Figure 12. Average classification accuracies for single ($q_2 \in \mathcal{Q}$) vs. double compression ($q_1 \rightarrow q_2, q_1 \in \mathcal{Q}_{q_2}^\Delta \setminus \{100\}$) detection, $\Delta = 20$. Feature space dimension 60.

grained tampered JPEG image detection via DCT coefficient analysis,” *Pattern Recognition*, vol. 42, no. 11, pp. 2492–2501, 2009.

- [16] T. Bianchi and A. Piva, “Image forgery localization via block-grained analysis of JPEG artifacts,” *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 3, pp. 1003–1017, 2012.
- [17] F. Pérez-González, G. L. Heileman, and C. T. Abdallah, “Benford’s law in image processing,” in *IEEE International Conference on Image Processing*, vol. 1, 2007, pp. 405–408.
- [18] D. Fu, Y. Q. Shi, and W. Su, “A generalized Benford’s law for JPEG coefficients and its applications in image forensics,” in *Security and Watermarking of Multimedia Content IX*, ser. Proceedings of SPIE, E. J. Delp and P. W. Wong, Eds., vol. 6505, 2007, 65051L.
- [19] X. Chu, Y. Chen, M. Stamm, and K. J. R. Liu, “Information theoretical limit of media forensics: The forensicability,” *IEEE Transactions on Information Forensics and Security*, in press.
- [20] M. Kirchner and S. Chakraborty, “A second look at first significant digit histogram restoration,” in *IEEE International Workshop on Information Forensics and Security (WIFS)*, 2015.
- [21] F. Benford, “The law of anomalous numbers,” *Proceedings of the American Philosophical Society*, vol. 78, no. 4, pp. 551–572, 1938.
- [22] A. Berger, “A basic theory of Benford’s law,” *Probability Surveys*, vol. 8, pp. 1–126, 2011.
- [23] P. Diaconis, “The distribution of leading digits and uniform distribution mod 1,” *The Annals of Probability*, vol. 5, no. 1, pp. 72–81, 1977.
- [24] B. Li, Y. Q. Shi, and J. Huang, “Detecting doubly compressed JPEG images by using mode based first digit features,” in *IEEE Workshop on Multimedia Signal Processing*, 2008, pp. 730–735.
- [25] T. Pevný and J. Fridrich, “Detection of double-compression in JPEG images for applications in steganography,” *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 2, pp. 247–258, 2008.
- [26] G. Schaefer and M. Stich, “UCID – an uncompressed colour image database,” in *Storage and Retrieval Methods and Applications for Multimedia*, ser. Proceedings of SPIE, M. M. Yeung, R. W. Lienhart, and C.-S. Li, Eds., vol. 5307, 2004, pp. 472–480.
- [27] R. Böhme and M. Kirchner, “Counter-forensics: Attacking image forensics,” in *Digital Image Forensics: There is More to a Picture Than Meets the Eye*, H. T. Sencar and N. Memon, Eds. Springer, 2013, pp. 327–366.
- [28] C. Pasquini, P. Comesaña, F. Pérez-González, and G. Boato, “Transportation-theoretic image counterforensics to first significant digit histogram forensics,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014, pp. 2699–2703.
- [29] P. Comesaña and F. Pérez-González, “The optimal attack to histogram-based forensic detectors is simple(x),” in *IEEE International Workshop on Information Forensics and Security*, 2014, pp. 1730–1735.
- [30] V. Holub and J. Fridrich, “Digital image steganography using universal distortion,” in *ACM Workshop on Information Hiding and Multimedia Security*, 2013, pp. 59–68.
- [31] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*. John Wiley & Sons, 2012.