Conditions for Effective Detection and Identification of Primary Quantization of Re-Quantized JPEG Images

Matthew James Sorell
School of Electrical and Electronic Engineering
University of Adelaide SA 5005
Australia
+61 8 8303 3226

matthew.sorell@adelaide.edu.au

ABSTRACT

The choice of Quantization Table in a JPEG image has previously been shown to be an effective discriminator of digital image cameras by manufacturer and model series. When a photograph is recompressed for transmission or storage, however, the image undergoes a secondary stage of quantization. It is possible, however to identify primary quantization artifacts in the image coefficients, provided that certain image and quantization conditions are met. This paper explores the conditions under which primary quantization coefficients can be identified, and hence can be used image source identification. Forensic applications include matching a small range of potential source cameras to an image.

Categories and Subject Descriptors

I.4.2 [**Image Processing and Computer Vision**]: Compression (Coding) – *Approximate methods*; I.5.1 Models: *Statistical*; K.4.1 [**Computers and Society**]: Public Policy Issues – *Abuse and crime involving computers*

General Terms

Measurement, Security

Keywords

JPEG, Forensics, re-compression, source identification, digital camera.

1. INTRODUCTION

In previous work [1], the author demonstrated that the choice of Quantization Table in the JPEG image compression algorithm used in a digital camera is highly dependent on the particular camera manufacturer, and to a lesser extent on the model series. A sample set of over 5000 digital photographs were used to extract 330 Quantization Tables from 27 different camera models from 10 brands, and it was shown that just 42 Tables were common to more than one camera model. After aggregating the results by camera model *series*, just 25 tables were found to be common across more than one manufacturer, and of this set, 19 of a possible 20 were common to the manufacturers Nikon and Olympus, suggesting a common source of JPEG encoding algorithm.

The Quantization Table is a useful source discriminator in cases where metadata (notably Exif [2] metadata) has been removed or is suspected of having been modified. Under certain conditions, the effect of the original quantization can also survived subsequent compression, such that it is at least possible to narrow down the range of potential source cameras of a recompressed image of interest. Further, if the quantization history of a set of images can be established, it is possible to collate image sets by that quantization history.

In [1], a multi-hypothesis test based on the 330 sample Quantization Tables was demonstrated, using a weighted sum of matched filters for each of the 64 quantization coefficients. This paper takes that analysis further, by examining the conditions under which detection of two-stage quantization is possible, using these results to establish the subset of plausible primary Quantization Tables.

Previous work on this problem includes [3], which provided an incomplete analysis based on single photographs from 300 candidate cameras; [5], which used a neural network approach for pattern matching; and [6], which focussed on a maximum-likelihood approach for overall Quantization Table estimation.

2. JPEG COMPRESSION

The JPEG standard is defined in [4] and the details of the standard are given in [7] for the interested reader. There are a number of modes of operation of the JPEG compression algorithm, but we consider only the progressive mode which is designed for lossy compression of continuous-toned images and is ubiquitously implemented in digital cameras and image editing software. The JPEG compression stages are introduced briefly here.

2.1 Image Compression

The JPEG compression algorithm takes as its input three color planes representing Red, Green and Blue light. These undergo a reversible color-space transform into Luminance (Y), Chrominance Red (Cr) and Chrominance Blue (Cb). The Y plane is identical to that used in black-and-white television and contains high resolution information which stimulates the many broad-spectrum rods in the eye. The Cb and Cr planes represented the "coloredness" of the image, stimulating the color-sensitive optical cones, which are far less sensitive in resolution. The chrominance

E-FORENSICS 2008, January 21-23, Adelaide, Australia

DOI 10.4108/e-forensics.2008.2764

planes are therefore downsampled, or averaged, to reduce the data in these planes by 50% or 75%.

The next stage is to consider the three planes independently and perform a spatial frequency transformation, in the form of a two-dimensional discrete cosine transformation (DCT), on "microblocks" of 8-by-8 pixels. This transform is also reversible, but importantly it separates the representation of the microblocks into broad contrast and fine detail. As the eye is sensitive to broad contrast and fine textures, but is much less sensitive to broad detail, the transformed coefficients can now be coarsely quantized. The high spatial frequencies are rounded off in quite large quanta, while low spatial frequences are typically rounded off in significantly smaller quanta. It is this stage which provides significant lossy compression, as the corresponding rounded numbers can now be stored very efficiently using conventional lossless compression techniques.

2.2 Quantization Tables

The rounding off of the DCT coefficients is determined by Quantization Tables which are specific to the implementation of the JPEG compression software used by a particular manufacturer. A particular camera will often use multiple tables, usually to meet some specified range of file size or quality level.

Different Quantization Tables are used for the luminance versus the chrominance planes. Up to four Quantization Tables are supported by the standard, but in practice it is common to use just two (using a common table for the chrominance planes) or three.

The JPEG standard [4] provides example Quantization Tables, but the work in [1] identified these examples only in thumbnail preview images. In practice, the Quantization Table is an effective discriminator of manufacturer and camera model series.

3. DETECTION ALGORITHM

It is common practice for a photographic image to undergo recompression before being stored, transmitted, or uploaded to a website. Anecdotal evidence suggests that images of interest in this contextare rarely re-scaled, cropped or significantly enhanced, and that in most cases the transformation is nothing more than lossier compression of the image at the original resolution. Of course, this statement does *not* apply to images which have been deliberately tampered with, but the detection of such forgery is not within the scope of this paper.

When an image is re-compressed, the image is extracted into its original Red-Green-Blue representation (some software also works directly on the Y-Cb-Cr representations). It is very common for the chrominance planes to be further downsampled and for that reason we do not consider the chrominance planes any further in this context.

3.1 The Double Quantization Ripple Effect

Considering only the luminance plane, it should be apparent that if a coefficient is originally rounded off to say the nearest multiple of 2, and then requantized to the nearest multiple of three, the coefficients are no longer smoothly distributed. For example, the series 1-2-3-4-5-6-7-8-9-10 would be rounded off first to 2-2-4-4-6-6-8-8-10-10 and then to 3-3-3-3-6-6-9-9-9-9; whereas direct rounding would have resulted in the uniform 0-3-3-3-6-6-6-9-9-9. The result is a ripple effect in the distribution of coefficients, as shown in Figure 1. This ripple effect can be detected by matching with an appropriate template.

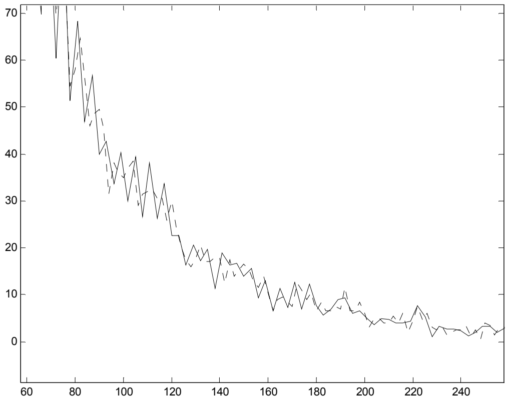

Figure 1: Double rounding leads to a periodic ripple in the distribution of coefficients. The dotted line shows the distribution of the coefficients in the original image, rounded off to 2. The solid line shows the distribution after rounding off to 3. The periodic behaviour is starkly evident.

However, this effect will not be detected under two key circumstances:

1. If Q_2 is divisible by Q_1, there will be no periodic ripple as the requantization is distributed evenly, or
2. If Q_2 is sufficiently large, the period of the ripple exceeds the range of the DCT coefficients, which lie in the range -1024 to +1023 (except for the first element, representing average intensity, which is supported over twice this range). In this case, there will be no effect on which to base a template.

Figure 2 summarizes these two effects. The white area represents pairs of (Q_1, Q_2) which have a detectable ripple effect, whereas the ripple effect is either non-existent, or out of range, in the black areas.

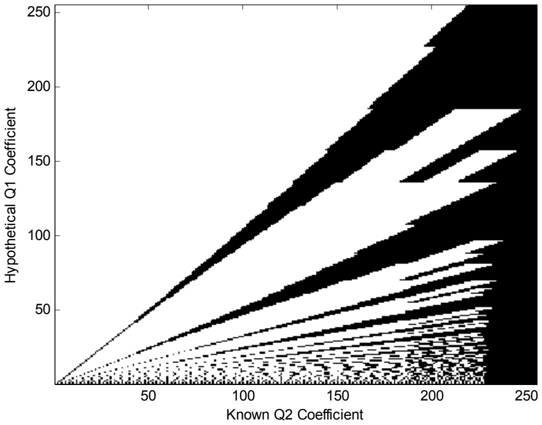

Figure 2: Showing the Quantization Coefficient pairs (Q_1, Q_2) which result in the ripple effect (white zone), versus areas where Q_1 cannot be deduced using the ripple effect (black zone).

3.2 Matched Filter Correlation

Our detection algorithm is based on a matched filter. The algorithm proceeds as follows:

1. For each j of the 64 elements of the macroblocks, compute a histogram of the *absolute* value of the jth coefficient over the entire image. This results in a single-sided generalised Gaussian distribution, potentially with the ripple effect intact as shown previously in Figure 1.
2. Apply a symmetric high-pass filter to remove the underlying distribution and retain the ripple effect:

$$\overline{x_k}(i) = \frac{2x_k(i) - x_{k-1}(i) - x_{k+1}(i)}{2}$$

   For reasons of symmetry, $x_0(i)$ is doubled so as to reduce the impact on $x_1(i)$. It should be noted that coefficients which are rounded to 0 do not contribute to a detectable quantization effect and so are removed from consideration. It should be noted that due to the Central Limit Theorem, the underlying "noise" in the signal, in fact artifacts of the source image, can be approximated as Additive White Gaussian Noise, and hence the use of a matched filter for detection is optimal.

3. A template is then created for a given pair of primary and secondary coefficients based on a uniform distribution of coefficients. Both the template and the filtered signal are normalised by the root-mean-square to ensure that their correlation is normalised to the range -1 to +1. In addition, the template for any pair in the "black region" of Figure 2 is set to a vector of zeros.
4. The correlation of the filtered signal and the template is then computed for each of the 64 elements of the macroblock:

$$c(Q_1(i), Q_2(i)) = \sum_{k=-1023}^{1023} \overline{x_k} \cdot \overline{t_k}(Q_1(i), Q_2(i))$$

3.3 Decision Rule

In [1], the correlation coefficients were scaled by the number of non-zero elements in the original histogram and then summed to provide a weighted score against the hypothetical original Q_1 Quantization Table. That approach was based on some experimental results indicating effectiveness, but was not intended to be optimal. Rather, it intended to demonstrate the feasibility of identifying the primary Quantization Table.

The approach presented here is quite different, and we are not attempting to identify the one hypothesis which best matches the data in some sense. Rather, it is our intention here to consider each of the 64 correlation coefficients and hence minimise the size of the set of feasible primary Quantization Tables. Further analysis can then be undertaken on the reduced set, if required.

We consider three outcomes from the correlation test:

1. The correlation test is positive, indicating a possible match. The greater the correlation value, the better the match
2. The correlation test is zero, indicating either that we are operating in the "black zone" where there is no ripple effect, or that there are no non-zero elements in the histogram for that particular histogram, or that the specific data cancels itself precisely
3. The correlation test is negative, indicating a possible mismatch.

It should be noted that the characteristics of the underlying photograph act as noise in the context of these results, and that it is possible through the specific textures of an image to obtain a strong match, or mismatch, through coincidence.

A number of variations on this technique were reviewed by experiment using the 13 secondary Quantization Tables of Adobe Photoshop and 10 secondary Quantization Tables of Microsoft Photo Editor, corresponding to quality levels 10%, 20%, ..., 100%. It should be noted that a quality level of 1% in Photo Editor results in all Quantization Coefficients being equal to 255, and therefore operates in the "black zone". A total of 330 distinct primary Quantization Tables were considered from 4487 photographs.

The variations we considered included:

1. Considering positive correlations to be detected above some positive non-zero threshold, and similarly considering negative correlations to be detected below some negative non-zero threshold. In practice, it was found that this had very little impact on the results, and so the simpler positive-zero-negative approach is used.
2. Taking into account whether a zero result is due to the double quantization "black zone" or some other reason. Importantly, it was found that this approach was quite critical in analysing results. Hence, the *number of valid tests* is defined as the number of coefficient pairs (Q_1, Q_2) in the "white zone", where Q_1 is defined as the true underlying primary quantization coefficient, provided that there are non-zero samples from which to derive a signal (see Step 2 in Section 3.2).
3. Considering whether to simply consider the number of positive correlations, or whether to subtract the number of negative correlations. Experimental results showed better performance if the positive results were weighed up against the negative results. Hence, a *primary Quantization Table hypothesis* is considered *feasible* if the number of positive correlations, less the number of negative correlations, is positive. Otherwise the hypothesis is *unsupported*.

4. ANALYSIS OF RESULTS

Our results are analysed according to the number of valid tests as defined in section 3.3. In [1], results were categorised by the maximum coefficient in the true primary Quantization Table, but this was found not to be a particularly effective technique for establishing whether a specific primary Quantization Table would be easily detected against a specific secondary Quantization Table. In contrast the *number of valid tests* approach is an objective indicator of the quality of the data presented to the decision rule.

We considered through indicators of performance in our results:

1. The probability that the true primary Quantization Table would be included in the set of *feasible* hypotheses,
2. The probability that the true primary Quantization Table would be the highest-ranked result, that is that there would be no other hypothesis with a high net number of positive correlations. As the decision rule is based on integers, it

should be noted that this approach allows for several equal first-ranked hypotheses.

3. Given that the true primary Quantization Table is considered feasible, its mean ranking by net number of positive correlations

4.1 Secondary Quantization by Adobe Photoshop

All of the 13 Quantization Tables implemented by Adobe Photoshop were considered. Three representative examples are given here.

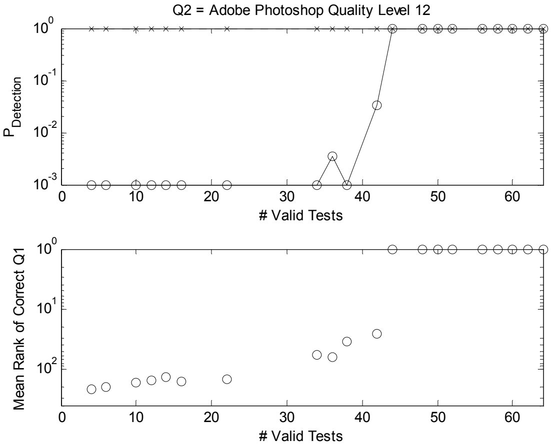

Figure 3: Results for Adobe Photoshop using Quality Level 12. The upper plot shows the number of first-ranked detections of the true primary Quantization Table (solid line with circled points) and the probability of the true primary Quantization Table being considered *feasible* (dotted line with crosses). The lower graph shows the mean rank of the true primary Quantization Table if it is in the feasible set. Note that the y-axis is inverted so that a rank of "1" is at the top of the axes. The x-axis counts the number of valid hypotheses, that is the number of quantization pairs in the "white zone" of Figure 2.

Figure 3 shows the results for the highest-quality setting for Photoshop. To ensure statistical validity, data points are aggregated in steps of two valid tests and are only shown if there are more than 10 samples. It can be seen that where there are more than 44 valid tests, the true primary Quantization Table is always ranked first in the decision rule.

Figure 4 shows the results for Photoshop Quality Level 10. Although this high-quality setting results in a high probability of capturing the true primary Quantization Table in the feasible hypothesis set, it is clear that the true primary Quantization Table no longer dominates the ranking. This same behaviour is seen as the Quality Level is reduced below 6, noting that due to a bug in implementation, Quality Level 6 results in a higher-quality image than Quality Level 7.

Figure 5, showing Quality Level 4, shows surprisingly similar results to Quality Level 10, and in fact this is true even for Quality Level 0. The reason for this, perhaps unexpected, result, is that Photoshop limits its secondary quantization coefficients for high spatial frequencies to 15, while quantizing low spatial frequencies more coarsely at low qualities. This results in images which have particular poor visual quality at low quality levels compared with JPEG images of similar file size compressed using other software. It also means that high frequency distributions are largely retained by images recompressed using Photoshop, regardless of the Quality Level used.

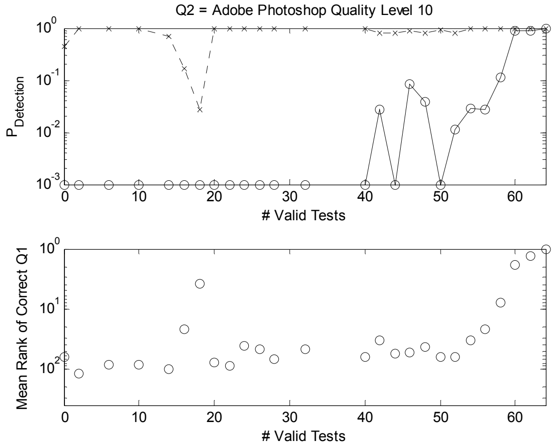

Figure 4: Results for Adobe Photoshop using Quality Level 10. See Figure 3 caption for details.

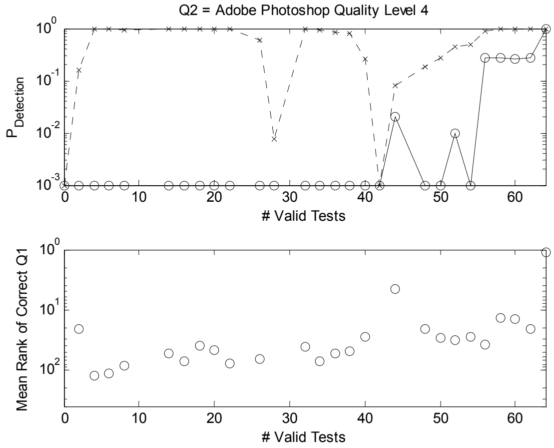

Figure 5 Results for Adobe Photoshop using Quality Level 4. See Figure 3 caption for details.

4.2 Secondary Quantization by Microsoft Photo Editor

The choice of Quantization Table for images compressed using Microsoft Photo Editor is quite different to Photoshop. It appears that a single nominal Quantization Table has been specified at some nominal quality level. As the Quality Level is increased towards 100%, the Quantization Table is scaled down uniformly, resulting in finer quantization. At 100%, all quantization coefficients are equal to one, meaning that there is no further rounding and that primary quantization is particularly simple to detect. On the other hand, at 1% all quantization coefficients are scaled up to their maximum value, 255, and there is no detectable secondary quantization effect.

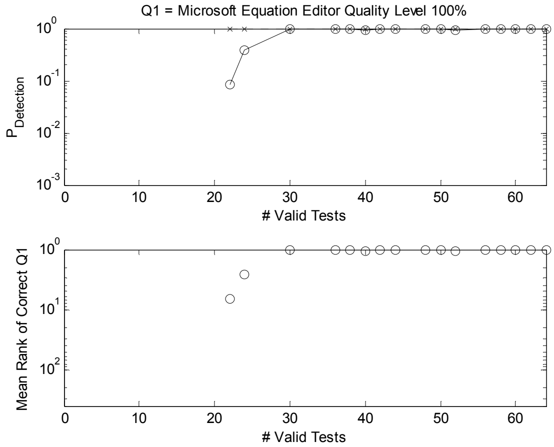

Figure 6: Results for Microsoft Photo Editor with Quality Level set to 100%. See Figure 3 caption for details.

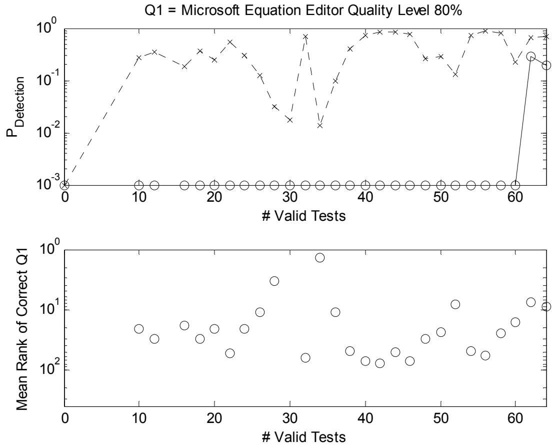

Figure 7: Results for Microsoft Photo Editor with Quality Level set to 80%. See Figure 3 caption for details.

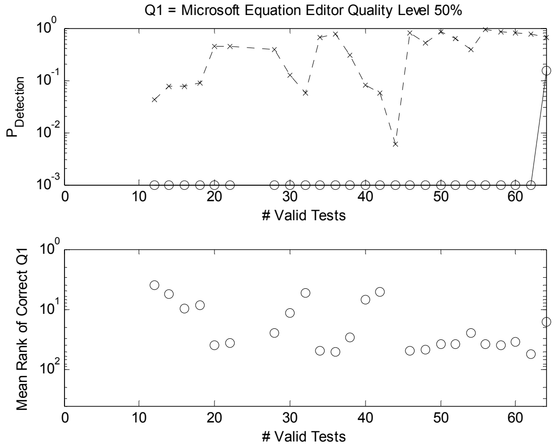

Figure 8: Results for Microsoft Photo Editor with Quality Level set to 50%. See Figure 3 caption for details.

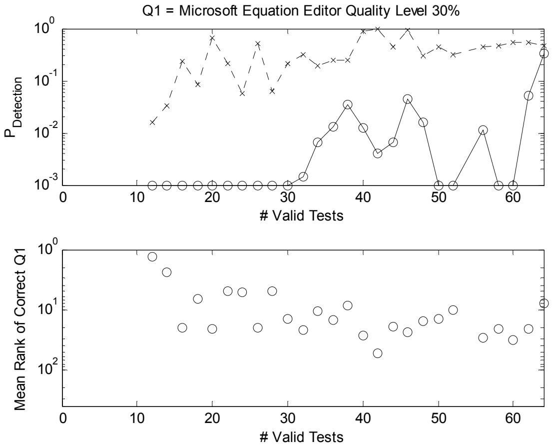

Figure 9: Results for Microsoft Photo Editor with Quality Level set to 30%. See Figure 3 caption for details.

Figure 6, showing results at Quality Level 100%, indicates very strong performance as expected. There is no secondary rounding and so as few as 32 valid tests results in a high rate of first-ranked true primary Quantization Table identification.

As the Quality Level is reduced, however, there is a sharp decline in performance. While the probability of inclusion of the true primary Quantization Table in the feasible set remains high, the ranking declines to an average of typically 10 to 100. This is still useful, as it allows for prioritizing certain camera types over others in subsequent forensic analysis, but the results are certainly not as clear-cut as those for images recompressed using Adobe Photoshop. Furthermore, the probability of the true hypothesis being first ranked rapidly drops away to zero. These results are summarised in Figures 7, 8 and 9.

The primary reason for this very different behavior is clear. Firstly, Photo Editor quantizes high frequencies more coarsely than lower frequencies, and maintains the ratio of the quantization coefficients regardless of the quality level. This means that the ripple effect rapidly becomes undetectable as the quality level decreases, in contrast to Photoshop, in which the high frequencies are retained.

Furthermore, as Photo Editor does not recognise Exif metadata and strips it from the recompressed file, it is of some concern that recompression using Photo Editor, which is shipped with the Windows operating system, can easily suppress forensic traces of the primary image source in a way that other software, most notably Adobe Photoshop, cannot.

5. CONCLUSIONS

This paper has emphasised a detection and estimation approach for identification of the primary quantization coefficients of a recompressed JPEG image based on determining which hypothetical primary Quantization Tables are feasible, rather than attempting to determine the single candidate with the best match. It has been demonstrated that the latter approach is not particularly effective and that the double-quantization ripple effect is better exploited to reduce the size of the candidate set of primary Quantization Tables.

The implications of this work are clear in forensic investigation. By identifying the primary quantization coefficients of an image, or as small a set of coefficients as possible, it is possible to narrow down the range of potential image source devices. When considered in conjunction with other evidence, such as unreliable metadata, image resolution, camera fingerprinting, physical evidence, etc, it is clear that the quality of the overall set of evidence is improved.

A further application, requiring some minor re-working of the analysis given here, is to consider which photographs from a large set have undergone the same quantization history. Q2 is explicitly given in every case, and by identifying images with matching primary quantization "fingerprints", it would be possible to segment a large set of images into smaller sets for human inspection.

6. ACKNOWLEDGMENTS

I would like to acknowledge the assistance of Jeremy Messé for his work in setting up the reference photograph database and coding many of the algorithms on which this paper is based.

7. REFERENCES

[1] Sorell, M. J. 2008. Digital Camera Source Identification Through JPEG Quantization. In Multimedia Forensics and Security", Chang-Tsun Li (ed), Idea Group, *in press*

[2] Japan Electronics and Information Technology Industries Association. 2002. JEITA CP-3451 Exchangeable image file format for digital still cameras: Exif Version 2.2, Japan Electronics and Information Technology Industries Association, April 2002.

[3] Farid, H, 2006. Digital Image Ballistics from JPEG Quantization, Dartmouth Computer Science Technical Report, TR2006-583, September 2006

[4] International Telecommunications Union, 1993. CCITT T.81 Information Technology – Digital Compression and Coding of Continuous-Tone Still Images – Requirements and Guidelines, International Telecommunications Union, 1993

[5] Lukas, J., and Fridrich, J. 2003. Estimation of Primary Quantization Matric in Double Compressed JPEG Images, Proc. of Digital Forensic Research Workshop, Cleveland, Ohio, August 2003.

[6] Neelamani, R., de Queiroz, R., Fan, Z., Dash, S., & Baraniuk, R., 2006. JPEG compression history estimation for color images, IEEE Trans. on Image Processing, Vol 15, No 6, pp. 1365-1378, June 2006

[7] Wallace, G.K. 1991. The JPEG Still Picture Compression Standard, Communications of the ACM, Vol. 34, No. 4, pp 30-44, April 1991.