Comparison of image-based functional monitoring through resampling and compression

Steven J. Simske, Margaret Sturgill, Jason S. Aronoff

HP Laboratories
HPL-2009-145

Keyword(s):
Image forensics, counterfeit detection, classification, accuracy, lossy compression, down-sampling

Abstract:

Image-based applications such as remote surveillance, environmental monitoring, and robotic navigation are often bandwidth-limited, and benefit from image down-sampling or compression. Often a decision is made without considering the relative impact on the functional goal of the monitoring of the different down-sampling and/or compression choices. In this paper, we use a specific "remote" monitoring application - the distinction between images of authentic products and counterfeit products - to assess the impact of down-sampling and compression on the classification accuracy of the counterfeit detection imaging software.

External Posting Date: June 21, 2009 [Fulltext] Approved for External Publication
Internal Posting Date: June 21, 2009 [Fulltext]

To be published in IEEE International Geoscience & Remote Sensing Symposium (Cape Town, South Africa)

COMPARISON OF IMAGE-BASED FUNCTIONAL MONITORING THROUGH RESAMPLING AND COMPRESSION

Steven J. Simske, Margaret Sturgill, Jason S. Aronoff

Hewlett-Packard Labs, 3404 E. Harmony Rd. MS 36, Fort Collins CO USA 80528

ABSTRACT

Image-based applications such as remote surveillance, environmental monitoring, and robotic navigation are often bandwidth-limited, and benefit from image down-sampling or compression. Often a decision is made without considering the relative impact on the functional goal of the monitoring of the different down-sampling and/or compression choices. In this paper, we use a specific "remote" monitoring application—the distinction between images of authentic products and counterfeit products—to assess the impact of down-sampling and compression on the classification accuracy of the counterfeit detection imaging software.

Index Terms— Image forensics, counterfeit detection, classification, accuracy, lossy compression, down-sampling

1. INTRODUCTION

The ubiquity of mobile cameras has made possible new imaging applications allowing consumers to interact with physical/printed materials in the environment, such as signage [1] and location-specific symbology [2]. Related applications include consumer interrogation of product packaging. Numerous organizations, including the Open Mobile Alliance [3] and the GS1 Mobile Com Extended Packaging project [4], also connect the consumer to the branded product through imaging-based services.

For security applications, bar codes and/or RFID chips are used to provide EPCglobal [5] mass serialization along with other security information (unique IDs, digital signatures of other printed information, etc.). These and other so-called security printed deterrents use valuable "real estate" (dedicated area to print) on the printed material, and so may conflict with the product branding and messaging. Example of security deterrents are given in Figure 1.

Many printed materials cannot accommodate security deterrents due to space (e.g. labels and medallions) or aesthetic (e.g. corporate/branded documents) concerns. As a consequence, we are interested in supporting, where possible, a deterrent-free approach. As mobile camera imaging capabilities continue to increase, this becomes feasible for more and more capture devices. In fact, it may soon be the case that network bandwidth, and not image quality, will be the primary consideration for mobile product authentication. To address this issue, we herein consider multiple image compression and resampling strategies to see if product authentication could be provided with significantly reduced bandwidth. Section 2 describes the experiments performed; Section 3 highlights the results; and we interpret these results in Section 4.

Figure 1. Samples of three different security printing deterrents: 2D barcode (left), color barcode with microtext (center) and 1D barcode (right).

2. EXPERIMENTS PERFORMED

We obtained original packages for 10 authentic and 10 counterfeit HP inkjet cartridges. Five different types of images were each scanned at 600x600 dots/inch (dpi) horizontal x vertical resolution using a desktop (HP Scanjet 8200), as shown in Figure 2. These are a set of two barcodes (hereafter "Barcode"), a blue spot color region ("Blue"), a set of color targets for print quality assurance ("Color"), a set of five branding images separated by whitespace ("Images") and a single large image ("Meadow"). The image areas were approximately 2.7, 2.8, 2.3, 4.2 and 10.1 in^2, respectively.

Ten image processing measurements, comprising the feature set, were computed for each of these 100 images (5 different images each from the 10 authentic and 10 counterfeit packages). Image entropy ("**e**", Equation 1) and standard deviation of the image histogram ("**Std(H_I)**", Equation 2), were computed from the intensities of the individual pixels, p_k, as shown below. Other metrics computed were the percent of pixels with largest relative neighborhood variance ("**%Edges**"), the mean value for these edges ("$\boldsymbol{\mu_{Edge}}$"), and the mean pixel variance, $\mathbf{PV_{xy}}$, based on the local differences in pixel intensity (Equation 3).

$$e = -\sum_{k=0}^{255} \frac{p_k}{\sum_{k=0}^{255} p_k} * \ln\left(\frac{p_k}{\sum_{k=0}^{255} p_k}\right) \qquad \textbf{Equation 1}$$

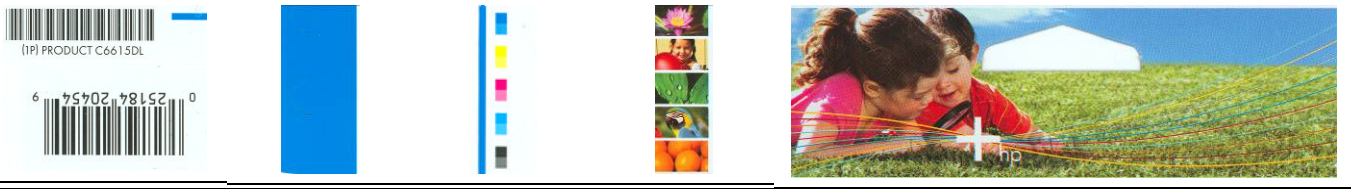

Barcode	Blue	Color	Images	Meadow
(1.8x1.5)	(1.1x2.5)	(0.9x2.5)	(1.2x3.5)	(4.4x2.3)

Figure 2. Samples of each of the five image types (actual image sizes in inch x inch at 600 dpi in parentheses)

$$\mathrm{Std}(H_I) = \frac{\sum_{k=0}^{255} p_k * (k-\mu)^2}{\sum_{k=0}^{255} p_k} \qquad \textbf{Equation 2}$$

$$PV_{xy} = \frac{(P_{x-1,y-1} - P_{xy}) + (P_{x-1,y+1} - P_{xy}) + (P_{x+1,y-1} - P_{xy}) + (P_{x+1,y+1} - P_{xy})}{4}$$

Equation 3

Mean saturation ("μ_{Sat}") was then computed, where saturation is defined as 255*(1-min(R,G,B))/(R+G+B). Mean connected component region size and variance were computed after thresholding the images based on pixel intensity ("$\mu_{Size\text{-}Intensity}$" and "$\mu_{Size\text{-}Intensity\text{-}\sigma*\sigma}$", respectively) and based on saturation ("$\mu_{Size\text{-}Saturation}$" and "$\mu_{Size\text{-}Saturation\text{-}\sigma*\sigma}$", respectively). These metrics were also computed for the same images down-sampled using ImageMagick [6] "–resample" option to 10 (Images, Meadow only), 20, 30, 40, 50, 60, 75, 100, 150, 200 and 300 dpi vertical x horizontal resolution; and for images compressed using JPEG to 1% or 2% original image size (effectively 60 and 85 dpi, respectively) before and after (effectively 30 and 42 dpi, respectively) down-sampling to 300 dpi. Down-sampling is abbreviated as DS, and Jpeg compression as JC. To indicate the effect of DS and JC on file size, we adopt the terminology DS_x = down-sampled by factor x; and JC_y = Jpeg-compressed by factor y. Thus, since DS is performed in both directions, downsampling to 10, 20, 30, 40, 50, 60, 75, 100, 150, 200 and 300 dpi is designated DS_{3600}, DS_{900}, DS_{400}, DS_{225}, DS_{144}, DS_{100}, DS_{64}, DS_{36}, DS_{16}, DS_9, and DS_4. Jpeg compression to 1% and 2% of file size are designated JC_{100} (to $1/100^{th}$ size) and JC_{50}, respectively. Since the Jpeg compression was preceded by down-sampling in two cases, these image transformations are referred to as (DS_4,JC_{100}) and (DS_4,JC_{50}).

3. RESULTS

The five different types of images illustrated in Figure 2 varied greatly by image metric. Seven of the metrics (**e**, **Std(H_I)**, **%Edges**, $\mu_{Size\text{-}Intensity}$, $\mu_{Size\text{-}Intensity\text{-}\sigma*\sigma}$, $\mu_{Size\text{-}Saturation}$, and $\mu_{Size\text{-}Saturation\text{-}\sigma*\sigma}$) distinguished the "Images" and "Meadow" classes from the other 3 classes, and $\mu_{Size\text{-}Saturation\text{-}\sigma*\sigma}$ distinguishes the class "Images" from the class "Meadow". The metrics $\mathbf{PV_{xy}}$ and μ_{Sat} provide assignment of the remaining images to the "Barcode", "Blue" and "Colors" classes. Thus, a decision tree was used to assign each original image to one of these five types based on one or more of these 10 "predictive" metrics. A previously described classifier [7] was then used to identify "authentic" and "counterfeit" images for each of these five image types. Each metric is assigned a critical point (**CPt**, see Figure 3) to one side of which it is assigned to "authentic" and the other to "counterfeit".

Table 1. Minimum down-sampling resolution (original image at 600 dpi in each direction) at which α_{CPt}, indicative of feature classification accuracy, was greater than of equal to that achieved for the original (600dpi) image.

Parameter	Barcode	Blue	Color	Images	Meadow
e	100	600	600	40	20
Std(H_I)	100	600	50	10	10
%Edges	600	20	200	20	10
μ_{Edge}	600	20	200	10	40
$\mathbf{PV_{xy}}$	600	600	600	40	20
μ_{Sat}	20	20	20	10	10
$\mu_{Size\text{-}Intensity}$	150	600	50	10	10
$\mu_{Size\text{-}Intensity\text{-}\sigma*\sigma}$	20	20	40	300	30
$\mu_{Size\text{-}Saturation}$	600	150	300	20	20
$\mu_{Size\text{-}Saturation\text{-}\sigma*\sigma}$	100	600	300	20	10

Previous work [8, 9, 10] has shown that statistical image metrics can predict image degradation and also be used to grade the quality of image restoration. We applied this approach herein to assess the functional metric of correctly assigning an image to "counterfeit" or "authentic". The

binary classifier [7] provides a comparative metric for accuracy, the statistical confidence at the critical point (α_{CPt}), which was used to determine the "functional monitoring" capability of the down-sampled images.

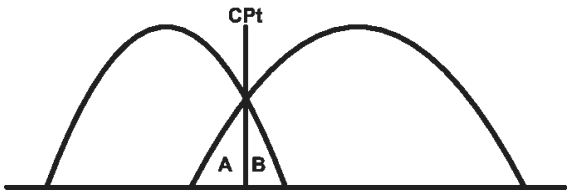

Figure 3. Critical point (CPt) between two populations is where area A = area B. Either of these areas is equal to the 1-α(CPt) as described in text.

That is, for each metric (Table 1), and for the mean of all 10 metrics (as described in [7]), we computed α_{CPt} and compared it to the α_{CPt} measured for the original 600 dpi images. If α_{CPt} for the down-sampled or compressed images >= α_{CPt} for the 600 dpi original images, then we are better off transmitting the smaller images.

Table 2. Original and transformed—D_x = down-sampled by factor x; and/or JC_y = Jpeg-compressed by factor y—images and the accuracy of classification. Original classification accuracy (top data row) is given in bold and italics. Any transformed image sets with higher classification accuracy than the original images are shown in boldface in the other rows.

Image Type	Barcode	Blue	Color	Images	Meadow
Original	***.896***	***.708***	***.788***	***.816***	***.743***
DS_4,JC_{100}	.643	.669	.719	.737	**.900**
DS_4,JC_{50}	.645	.668	.727	.737	**.900**
JC_{100}	.698	**.766**	.746	.740	**.884**
JC_{50}	.702	**.774**	.757	.740	**.892**
DS_4	.832	**.773**	**.846**	**.834**	**.801**
DS_9	.828	.674	**.798**	**.819**	**.880**
DS_{16}	.774	.675	.726	**.954**	**.961**
DS_{36}	.739	.682	.678	**.954**	**.893**
DS_{64}	.734	.669	.645	**.944**	**.836**
DS_{100}	.722	.636	.655	**.917**	**.882**
DS_{144}	.710	.648	.702	**.945**	**.880**
DS_{225}	.732	.642	.707	**.960**	**.865**
DS_{400}	.731	.668	.674	**.896**	**.834**
DS_{900}	.740	.662	.639	**.866**	**.768**
DS_{3600}	--	--	--	.718	**.806**

For the mean of all ten metrics, the α_{CPt} was .896, 0.708, 0.788, 0.816 and 0.743 (i.e. between 70%-90% accuracy) for Barcode, Blue, Color, Images and Meadow types, respectively (Table 2). The overall classification accuracy (using the binary classifier and all ten metrics) was 100% for all image sizes of the Barcode, Images and Meadow classes sizes; for the Color class image sizes from 30-600 dpi and when compressed from the 600 dpi original; and for the image sizes 40, 50, 75, 300 and 600 dpi for the Blue class. So, the classifier [7] was, in general, effective at differentiating authentic from counterfeit images.

We define $@A_{TF}$, or at-accuracy throughput factor, as the relative number of images (compared to 1 at 600 dpi) that can be successfully assigned to "counterfeit" or "authentic" classes with an accuracy >= the accuracy obtained for the original 600 dpi images, while using the same overall size in memory. For down-sampling, $@A_{TF}$ was 1, 4, 9, 900 and 3600, respectively, for the Barcode, Blue, Color, Images and Meadow image types. For compression, only the Blue and Meadow classes had a $@A_{TF}$ value greater than 1: specifically, 100 for the Blue images compressed from 600 dpi originals, and 400 for the Meadow images. That is, all four compression approaches for the Meadow images resulted in higher classification accuracy than for the uncompressed original images.

4. DISCUSSION AND CONCLUSIONS

Our results support the following approach to remote functional monitoring: (1) classify the image; (2) determine the smallest image (either down-sampled or compressed) for which α_{CPt} is >= to α_{CPt} for the 600 dpi original image; (3) transmit this smaller image; and (4) perform functional monitoring (in this case, correctly classifying each image). In this way, we achieved a higher overall accuracy by sending the Barcode images unaltered; compressing the Blue images by a factor of 100; and down-sampling the Color, Images and Meadow images by factors of 9, 900 and 3600, respectively. This overall system results in an effective $@A_{TF}$ of 4.45; that is, improved accuracy is achieved with a 4.45 reduction in transmission bandwidth.

The strategies described herein are examples of what we define as "functional imaging", wherein the transformations performed on the image are selected by the task or workflow to be completed. For three of the image types investigated, a substantial, lossy reduction in image size was possible without reducing the accuracy of classification. This significant increase in throughput of images for counterfeit detection is achieved without sacrificing accuracy. It is important to stress that the goal is not to transmit images that might be visually pleasing, but rather ones that contain enough information to perform the task (Figure 4). In contrast to previous studies [11, 12, 13] on compression and image classification, however, classification accuracy was actually shown to improve with increased image down-sampling and/or compression.

The results for classification are at first counterintuitive: smaller images actually classify with higher overall accuracy than the originals. This may be a consequence of the classifier used. The classifier selected [7] is designed to work best with Gaussian data, and the down-sampling operation—as well as many of the Jpeg

compression settings utilized, in which considerable loss of frequency information is obvious when viewing—is an averaging operation. The image metrics of the down-sampled and/or Jpeg compressed images, therefore, are likely more Gaussian than the metrics of the original images.

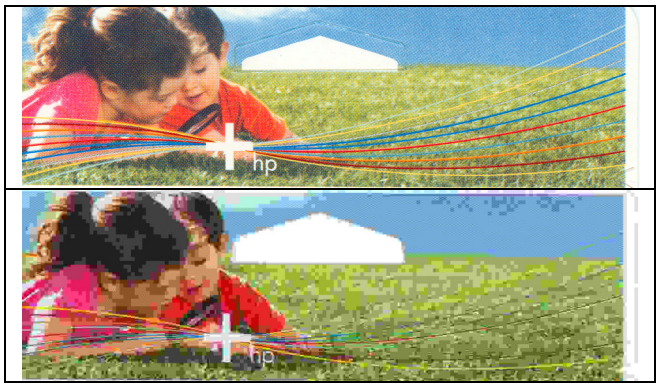

Figure 4. Meadow image at original 600 x 600 dpi (top) and the corresponding down-sampled and Jpeg compressed, (DS_4,JC_{100}) version of the Meadow image (bottom).

The smallest images created (DS_{14400} and DS_{3600}, for most of the images) were too small to assess because several of the features in the feature set—$\mu_{\text{Size-Intensity}}$, $\mu_{\text{Size-Intensity-}\sigma*\sigma}$, $\mu_{\text{Size-Saturation}}$ and $\mu_{\text{Size-Saturation-}\sigma*\sigma}$, in particular—were not calculable for these small images due to the small size of the regions formed. However, only the Meadow images showed increased classification accuracy at the smallest size investigated, so we did not alter the feature set to allow even smaller sizes.

The preliminary findings reported here are very promising. Future work will focus on increasing the sample sizes per image type and class, and on using additional image types, to determine the repeatability and breadth of application for the approach. Additional classifiers will also be tested to see if the successful functional imaging approach for counterfeit image classification described here is an idiosyncratic consequence of the classifier chosen.

The experiments performed here used a single resampling (down-sampling) approach, and a single compression (Jpeg) approach. Additional resampling (e.g. Hermite, cubic, etc.) and compression approaches will be considered in future experiments. Finally, we will explore the optimal deployment settings for the system, based on the range of DS and/or JC transformations that increase $@\mathbf{A}_{\mathbf{TF}}$.

In addition to remote surveillance and environmental monitoring, the techniques described here can be used to determine image down-sampling or compression recommendations for other "functional monitoring" product authentication/counterfeit detection. Since the classification of images is functional—that is, dependent on the task at hand—we believe a consideration of whether to down-sample or compress (and by what factor) should be an important part of any image-based monitoring application.

5. ACKNOWLEDGMENTS

The authors acknowledge and thank Dave Kellar and George Guillory for providing the authentic and counterfeit packages.

6. REFERENCES

[1] J. Yang, J. Gao, Y. Zhang, X. Chen, and A. Waibel, "An Automatic Sign Recognition and Translation System," *Proc. Workshop Perceptive User Interfaces*, pp. 1–8, 2001.

[2] E. Toye, R. Sharp, A. Madhavapeddy, and D. Scott, "Using Smart Phones to Access Site-Specific Services," *Pervasive Computing*, IEEE, 4:60–66, Jan.-March 2005.

[3] http://www.openmobilealliance.org/.

[4] http://www.gs1.org/productssolutions/mobile/.

[5] http://www.epcglobalinc.org/home.

[6]ImageMagick, http://www.imagemagick.org/script/index.php.

[7] S. J. Simske, "Low-resolution photo/drawing classification: metrics, method and archiving optimization," Proceedings IEEE ICIP, IEEE, Genoa, Italy, pp. 534-537, 2005.

[8] D. Li and S.J. Simske, "Atmospheric turbulence degraded image restoration by kurtosis minimization," IEEE Geoscience and Remote Sensing Letters, IEEE, Vol. 6, No. 2, pp. 244-247, 2009.

[9] D. Li, R.M. Mersereau and S.J. Simske, "Atmospheric turbulence degraded image restoration using principal components analysis," IEEE Geoscience and Remote Sensing Letters, IEEE, Vol. 4, No. 3, pp. 340-344, 2007.

[10] D. Li, R.M. Mersereau and S.J. Simske, "Blur identification based on kurtosis minimization," Proceedings IEEE ICIP, IEEE, Genoa, Italy, pp. 905-908, 2005.

[11] I. Blanes, A. Zabala, G. Moré, X. Pons, and J. Serra-Sagristá, "Classification of hyperspectral images compressed through 3D-JPEG2000," in I. Lovrek, R.J. Howlett, and L.C. Jain (eds.): KES 2008, Part III, LNAI 5179, Berlin: Springer-Verlag, pp. 416-423, 2008.

[12] A. Zabala, X. Pons, R. Díaz-Delgado, F. García, F. Aulí-Llinás, and J. Serra-Sagristá, "Effects of JPEG and JPEG2000 lossy compression on remote sensing image classification for mapping crops and forest areas," Proceedings IGARSS, IEEE, pp. 790-793, 2006.

[13] F. Tintrup, F. DeNatale, and D. Giusto, "Compression algorithms for classification of remotely sensed images," Proceedings ICASSP, IEEE, pp. 2565-2568, 1998.