

# Detecting Content Adaptive Scaling of Images for Forensic Applications

Claude Fillion<sup>1,2</sup>, Gaurav Sharma<sup>1,3</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY

<sup>2</sup>Xerox Research Center Webster, Xerox Corporation, Webster, NY

<sup>3</sup>Department of Biostatistics and Computational Biology, University of Rochester, Rochester, NY

## ABSTRACT

Content-aware resizing methods have recently been developed, among which, seam-carving has achieved the most widespread use. Seam-carving's versatility enables deliberate object removal and benign image resizing, in which perceptually important content is preserved. Both types of modifications compromise the utility and validity of the modified images as evidence in legal and journalistic applications. It is therefore desirable that image forensic techniques detect the presence of seam-carving. In this paper we address detection of seam-carving for forensic purposes. As in other forensic applications, we pose the problem of seam-carving detection as the problem of classifying a test image in either of two classes: a) seam-carved or b) non-seam-carved. We adopt a pattern recognition approach in which a set of features is extracted from the test image and then a Support Vector Machine based classifier, trained over a set of images, is utilized to estimate which of the two classes the test image lies in. Based on our study of the seam-carving algorithm, we propose a set of intuitively motivated features for the detection of seam-carving. Our methodology for detection of seam-carving is then evaluated over a test database of images. We demonstrate that the proposed method provides the capability for detecting seam-carving with high accuracy. For images which have been reduced 30% by benign seam-carving, our method provides a classification accuracy of 91%.

**Keywords:** Content-aware resizing, seam-carving, image forensics, detection

## 1. INTRODUCTION

A byproduct of the digital age is the advent of sophisticated image editing tools, such as GIMP<sup>1</sup> and Adobe Photoshop<sup>TM,2</sup>, which allow digital images to be easily manipulated using a myriad of methods. If these images are to be used by our judicial system or the news media, it is important that their fidelity to the original scene be verifiable. Digital watermarking provides such a means of authentication<sup>3,4</sup>, but requires that the watermark be inserted at the time of recording. Although this can be accomplished with specially equipped digital cameras, the use of these cameras is very limited and for a vast majority of digital images veracity cannot be established by watermarking technology. Consequently detection of digital image manipulations has emerged as a significant area of interest for forensic applications<sup>5</sup>. Previous work in this area<sup>6,7,8,9</sup>, has contributed to methods for detecting various image manipulation techniques such as cloning, splicing, and re-sampling. A clever new method for resizing images called Content-Aware Image Resizing, or Seam-carving<sup>10</sup>, poses similar forensic challenges which have not been studied thus far. In this paper we propose and develop a method for detecting seam-carving of images.

Seam-carving is a versatile tool that can be used both for deliberate removal of objects and benign image reduction, in which perceptually important content is preserved. Both of these uses of seam-carving have implications from a forensic standpoint. Deliberate removal of objects is a common forensic challenge. The content dependent nature of benign image reduction via seam-carving also poses a problem for applications such as photogrammetry, journalism, and law because the process often introduces significant distortions in the geometrical relationships between objects within an image.

---

C. Fillion: Email: [claudio.fillion@xerox.com](mailto:claudio.fillion@xerox.com), Phone: 1-585-422-0241. G. Sharma: Email: [gaurav.sharma@rochester.edu](mailto:gaurav.sharma@rochester.edu), Phone: 1-585-275-7313. This work is supported in part by the Air Force Office of Scientific Research (AFOSR) under grant number FA9550-07-1-0017.

The manner in which seam-carving manipulates images is quite distinct from other methods used in creating digital forgeries. Cloning involves cutting and pasting a portion of an image to conceal a person or object. These duplicated regions have been shown to be detectable<sup>9</sup>. Splicing involves combining two or more images into a single document. It has been shown that this splicing disrupts higher-order Fourier statistics in a detectable manner<sup>11</sup>. Re-sampling involves resizing, rotating, or stretching portions of an image and requires re-sampling of the original image onto a new sampling lattice. Re-sampling introduces periodic correlations in the modified image which are detectable<sup>7</sup>. Seam-carving, on the other hand, does not have much in common with these existing techniques and therefore the detectability of seam-carving merits independent study. The fact that the seam-carving method for content-aware resizing is now implemented in Adobe Photoshop™ CS4 and as a plug-in (Liquid Rescale<sup>12</sup>) for GIMP means that seam-carved images will proliferate, and further justifies the need for our analysis.

Before detailing our study, let us give a brief overview of seam-carving.

### 1.1 Seam-Carving Background

Seam-carving is one of several recently developed content-aware image resizing methods<sup>13, 14</sup> and has gained a measure of popularity due to its ability to overcome the limitations of traditional scaling and cropping. Its content-aware behavior resizes an image based upon its content, whereas traditional methods frequently de-emphasize (isomorphic scaling), distort (anamorphic scaling) or remove (cropping) content that may be important to the viewer.

The general approach to content-aware resizing involves first identifying regions of interest (ROI) within an image and then removing non-ROI portions of the image. One of the clever aspects of the seam-carving approach is that it does not require ROI to be reduced to a few clearly defined areas. It instead generates an image importance map which defines the “ROI” on a pixel by pixel basis. Seam-carving then resizes the image by adding or removing connected pixel paths, or seams, that have the lowest accumulated energy. Multiple image importance (energy functions) can be used, such as a saliency map, entropy, and gradient. In particular, among these options, the gradient operator is a simple yet effective operator for determining image complexity, which we shall use for our description.

A seam is defined as an 8-connected path of low energy pixels crossing the image from top to bottom, or from left to right. A dynamic programming technique is used to select the optimal seams in each direction, which are defined as the seams with the lowest accumulated “energy” indicative of the combined importance of the pixels on the seam. For image reduction, seam selection ensures that, while preserving the image structure, more low energy pixels are removed and more high energy pixels are maintained.

Figures 1 and 2 provide a visual demonstration of a horizontal and vertical seam for a typical image, ucid00220, from the UCID database<sup>15</sup>.




Figure 1: Demonstration of the first horizontal seam to be removed: (a) original image; (b) L1-norm gradient map with seam; (c) cumulative energy map with seam; (d) original with seam.




Figure 2: Demonstration of the first vertical seam to be removed: (a) original image; (b) L1-norm gradient map with seam; (c) cumulative energy map with seam; (d) original with seam.

Seam selection and removal can also be purposely directed by assigning higher or lower weights to a certain portion of an image. Assigning higher weight increases the importance of that portion of an image and therefore directs seams away from that portion. Assigning negative (lower) weight to a portion of an image reduces its importance and therefore attracts seams through that portion. This allows forgeries to be created in which people or objects are removed from an image by seam-carving.

Seam-carving can be used to enlarge images as well by inserting seams in the optimal location. The seam value typically is obtained by averaging the value of the pixels on either side of the inserted seam.

## 2. SEAM-CARVING DETECTION

In order to detect seam-carving, we employ a pattern recognition/machine learning methodology, shown in Figure 3 in its generic form. An image is classified as seam-carved or non-seam-carved by means of a classifier that uses, as its inputs, a set of features computed from the image. The classifier is trained using a corpus of representative images that includes both seam-carved and non-seam-carved images. The classifier training data consists of the set of features and a classification label that indicates the seam-carved or non-seam-carved status for each of the images in the training corpus.

In our work, we utilize a support vector machine<sup>16</sup> (SVM) based classifier, which among current practically accessible options typically offers the best classification performance.

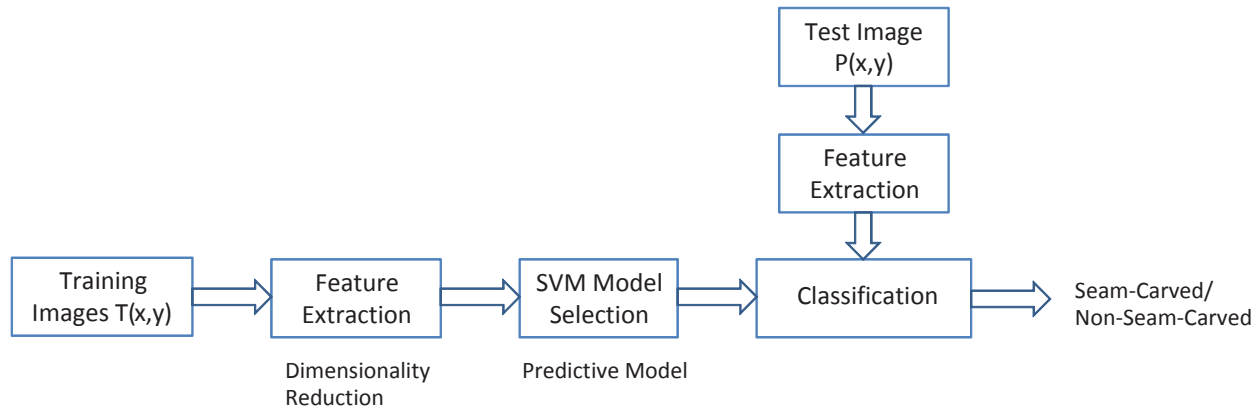


Figure 3: Block diagram of the machine learning methodology used for detecting seam-carving.

A classifier based approach presupposes the existence of meaningful features that enable detection of seam-carving. We have developed four sets of features for detection. Two of these feature sets are motivated by intuition based upon the seam-carving process, the third set is based upon the more general intuition that seam-carving will affect image statistics. Our fourth set of features is generated by applying an additional seam-carving operation to the image. This in turn is based on the assumption that a second seam-carving operation may have a less noticeable effect on our features because the image has already been disrupted by the first seam-carving operation – an approach that has shown some promise in closely related steganalysis applications<sup>17</sup>.

### 2.1 Features Based Upon Energy Bias

As mentioned previously, seam-carving affects the energy distribution within an image. Specifically, as a result of seam-carving, low-energy pixels are removed and a larger fraction of high energy pixels remain. Figure 4 demonstrates this bias for a typical image with 30% reduction. Figure 4(a) displays the cumulative distribution function (CDF) for a given image and is calculated as

$$CDF(i) = \frac{\sum_{j=1}^i \rho_j}{\rho_{tot}}, \quad (1)$$

where  $\rho_j$  is the pixel count for bin  $j$  of the gradient histogram and  $\rho_{tot}$  is the total number of pixels in the image.

Figure 4(b) displays the histogram for a typical image and demonstrates, in a non-cumulative manner, the effect of seam-carving across all gradient values. We also include results from the same image scaled by more conventional techniques – bicubic interpolation in particular, to demonstrate that, from an energy bias standpoint, scaled images typically behave similarly to the original image.




Figure 4: Energy distribution of an image: (a) 1024 bin histogram of gradient values; (b) cumulative distribution of histogram values. Shown on each plot is the original image (black), the image scaled by 30% (green), the image seam-carved by 30% (blue), and the image seam-carved by an additional 30% (red).

Based on these known behaviors we propose a set of features which analyze the distribution of gradient values for an image. We begin by selecting a feature that measures the percent of pixels within the image that fall below a predetermined energy threshold. Specifically we define this feature as

$$e_{low} = \sum_{i=0}^{\kappa-1} \frac{\rho_i}{\rho_{tot}}, \quad (2)$$

where  $\rho_i$  is the pixel count for bin  $i$  of the gradient histogram,  $\rho_{tot}$  is again the total number of pixels in the image, and  $\kappa$  denotes the number of histogram bins which fall below our energy threshold. Our current low energy threshold is 0.01.

Our second set of histogram-based features attempts to capture the mid-level histogram values. We accomplish this by first applying a zero-phase, low-pass filter to the gradient histogram plot in Figure 4(b) and then selecting sample points along the curve. We currently define mid-level histogram values as those having a normalized gradient value between 0.05 and 0.6.

Our third set of histogram-based features is directed at the information contained in the higher energy gradient values. In order to focus on the actual variation of higher gradient values, we first approximate removing the mean for each histogram by subtracting the low-pass filtered version of the histogram itself. We then capture the variation in the residual data in the form of a feature by taking the variance of the residual data over the range of high (greater than 0.6) gradient values.

Finally, we derive a group of histogram-based features from the CDF of Figure 4(a). While the CDF itself clearly shows energy bias, further separation (and therefore improved discrimination) is achieved by “normalization” of the CDF by  $e_{low}$  of (2),

$$nCdf(i) = \frac{\sum_{j=1}^i \rho_j}{\rho_{tot} e_{low}}, \quad (3)$$

as shown in Figure 5. We take a set of samples along the “normalized” CDF for our set of features.




Figure 5: Energy Cumulative Distribution normalized by  $e_{low}$

## 2.2 Features Based Upon Seam Behavior

Our next set of features hypothesizes that seam behaviors - such as distance between seams and energy along the path of a seam - are likely to be affected by the seam-carving process. Figures 6 and 7 demonstrate these seam metrics for a typical image. Figure 6(a)-(d) shows the path of the optimal seam at several points along the horizontal axis of the image before and after seam-carving by 30%. Thus in Figure 6(a) the magenta line traces the optimal path through 10% of the image, the cyan line traces the optimal path through 25% of the image, and so on. Figure 7 shows the cumulative energy of these paths at the same points. From these Figures we observe that the optimal seam path at given points across an image tends to be more dispersed after seam-carving. This behavior is somewhat common among seam-carved images and makes intuitive sense. Images which have not been seam-carved will tend to have more regions of uniform energy and thus the optimal seam path will be able to maintain its original path, or will likely find a minimal path nearby. As the image is seam-carved (and low energy regions are removed), the optimal seam path will be more likely to change as more of the image is traversed, and minimal paths will tend to be further apart. In terms of seam energy, this behavior will be manifested as a linear increase in cumulative seam energy across the image for an image which has not been seam-carved, and as a more disjointed set of increases after seam-carving.




Figure 6: Optimal seam path through 10% (magenta), 25% (cyan), 50% (green), 75% (red), and 100% (blue) of an image: (a) image before seam-carving; (b) cumulative energy map before seam-carving; (c) image after seam-carving; (d) cumulative energy map after seam-carving.

Based upon this intuition regarding seam behavior, we propose as features the distance between optimal seams at specified points along the image, and the cumulative energy at those points. For the horizontal seam in our example, we define the seam distance as

$$D_p = \sum_{j=1}^{N_p} \frac{|y_p(j) - y_{tot}(j)|}{N_p}, \quad (4)$$

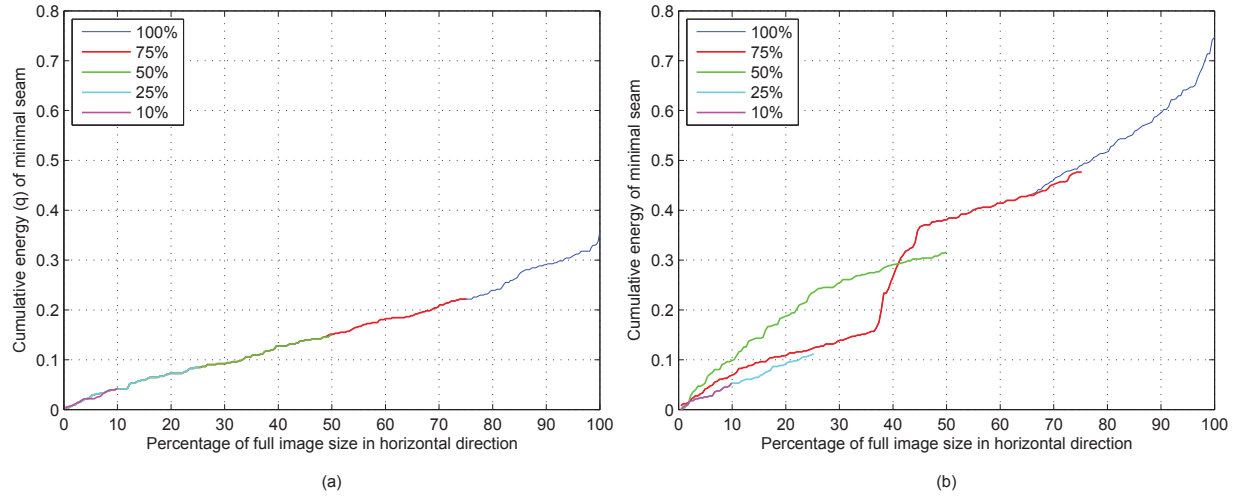


Figure 7: Cumulative seam energy of the original (a) and seam-carved (b) image

where  $D_p$  is the distance between the optimal seam through  $p$  percent of the image and  $p$  percent of the optimal seam through the entire image,  $N_p$  is the number of pixels through  $p$  percent of the image, and  $y_p$  and  $y_{tot}$  are the vertical coordinates of the respective seams. Vertical seam distance is calculated in a similar fashion.

The cumulative energy features are simply the cumulative energy at  $p$  percent normalized by the cumulative energy of the entire seam.

### 2.3 Features Based Upon Higher Order Image Statistics: Wavelet Absolute Moments

Our third set of features is predicated on the assumption that seam-carving will affect the statistics of the image. In this context, wavelet absolute moment features have proven to be effective in steganalysis<sup>18</sup>, and we examine these in our work as well. The features developed in<sup>18</sup> were calculated in the wavelet domain as higher-order absolute moments of the noise residual. In our work there is no known model to represent seam-carving, we therefore calculate the wavelet absolute moments of the image itself. We currently determine the optimal number of moments heuristically. We have evaluated our performance over a range of moments. In their steganalysis work Farid and Lyu<sup>19</sup> utilized 4 moments. Goljan *et al.* used 9 moments in<sup>18</sup> and Holotyak *et al.* used 33 moments in<sup>20</sup>. We have also examined the effect wavelet decomposition level upon performance.

We use the following normalized moment formulation for each of the three sub-bands and each wavelet decomposition level:

$$m_{H_p} = \begin{cases} \frac{\sum_i \sum_j |w_H(i, j) - \mu(w_H(i, j))|^p}{\sigma^p} & p \neq 2 \\ \frac{\sum_i \sum_j |w_H(i, j) - \mu(w_H(i, j))|^p}{ij} & p = 2 \end{cases} \quad (5)$$

where  $w_H$  are the wavelet coefficients of the horizontal sub-band of the wavelet transform and  $p$  is the moment number. The precise number of wavelet based features used by our method is  $n_{level} \times n_{subband} \times n_{moments}$ .

### 2.4 Features Based on Second Seam Carve

As mentioned earlier, our fourth class of features is derived from applying an additional seam-carving operation to the image. All of the previously mentioned features are again calculated.



### 3. EXPERIMENTAL RESULTS

There are three modes of seam-carving that we wish to detect. Benign image reduction, benign image enlargement, and directed (or deliberate) image reduction. Our framework is applicable to all three types of seam-carving that we wish to detect. We have tested our method over a set of test images consisting of 1484 uncompressed images, taken from a variety of sources. The images have been converted to grayscale and resized (using bicubic interpolation) in order to obtain additional image samples and to prevent our classifier from making predictions based upon image size. The sizes of the images in our experiments range from 307 x 205 to 512 x 683 pixels. For the benign reduction and enlargement cases, we partition our image set randomly into two halves, one of which is used for training and the other for testing. For each set of original images, a corresponding set of seam-carved images has been generated. We explore seam-carving based reduction ranging from 5% to 30%. For removal of selected objects, manual construction of the seam-carved images has been performed at a single resizing amount of 20%, and a smaller set of images (25 each for training and prediction) is tested.

Our test results demonstrate the ability to detect images that have been seam-carved using Adobe Photoshop™ CS4, based on the assumption that this is the most likely source of seam-carved images. We have developed a MATLAB™-based seam-carving implementation which enabled detailed study of seam-carving and allowed us to hypothesize the set of features based upon seam behavior. As expected, our ability to detect seam-carving using our own implementation was slightly better (by about 2%).

To date our work has focused primarily on benign reduction, but results are also included for benign enlargement and directed reduction.

#### 3.1 SVM Configuration

For our classifier, we utilized the LIBSVM<sup>21</sup> Support Vector Machine library. Parameter selection is important for selecting good SVM models. We employed the following set of parameters which have been shown to perform well in a variety of scenarios<sup>22</sup>: a soft margin C-SVC SVM classifier, a radial basis function kernel, and a 5-fold grid-search cross-validation process for selecting error cost and kernel gamma parameters.

#### 3.2 Detection of Benign Reduction

We begin by demonstrating the effectiveness of features based on energy bias, as described in Section 2.1. We define 46 features to capture the energy bias (histogram segments and normalized CDF) information of interest. Results for each of the four types of energy bias feature are shown in Table 1. From the table it is evident that each of the four types contributes to classification accuracy.

Table 1: Classification Accuracy Using Features Related to Energy Bias for a 30% Benign Reduction by Seam-Carving

Features	Classification Accuracy (percent)
Low Energy ( < 0.01)	71.77
Mid Energy ( 0.05 to 0.6)	70.69
High Energy ( > 0.6)	71.43
Normalized CDF	67.86
Combined	77.90

Next we demonstrate the effectiveness of features based on seam behavior, as described in Section 2.2. We define 18 features to capture the seam behavior of interest. Results for each of the four types of seam behavior features are shown in Table 2. The seam energy features appear to provide additional discrimination capability.

Table 2: Classification Accuracy Using Features Related to Seam Behavior for a 30% Benign Reduction by Seam-Carving

Features	Classification Accuracy (percent)
Horizontal Seam Distance	63.34
Horizontal Seam Energy	71.09
Vertical Seam Distance	63.01
Vertical Seam Energy	70.08
Combined	78.30

Table 3 shows results at several reduction amounts for a first seam-carve operation, a second seam-carve, and the combined effect of both seam-carving operations. For both the energy bias and seam behavior feature sets, the first seam-carving operation provides more discrimination capability but the second seam-carving operation also makes a meaningful contribution to the combined score.

Table 3: Classification Accuracy of Features Related to Energy Bias and Seam Behavior

Features	Classification Accuracy for Benign Seam-Carving of:			
	5%	10%	20%	30%
Energy Bias (1 <sup>st</sup> Seam Carve)	54.51	60.31	68.67	77.90
Energy Bias (2 <sup>nd</sup> Seam Carve)	53.77	58.96	65.97	74.73
Energy Bias (1 <sup>st</sup> and 2 <sup>nd</sup> Seam Carve)	57.21	64.49	74.73	82.41
Seam Behavior (1 <sup>st</sup> Seam Carve)	60.24	66.37	72.71	78.30
Seam Behavior (2 <sup>nd</sup> Seam Carve)	55.53	59.10	65.30	69.27
Seam Behavior (1 <sup>st</sup> and 2 <sup>nd</sup> Seam Carve)	58.89	68.80	78.03	83.89

We next turn our attention to features based upon higher order image statistics. We examined accuracy over a range of wavelet absolute moments and decomposition levels. For our set of training data, the best performance occurred when using 4 moments and 6 decomposition levels, which requires  $n_{\text{level}} \times n_{\text{subband}} \times n_{\text{moments}} = 72$  features.

Table 4 shows the performance of the Wavelet absolute moment features. As with our other features, the first seam-carving operation provides more discrimination capability but the second seam-carving operation also makes a meaningful contribution to the combined score.



Table 4: Classification Accuracy of Features Related to Higher Order Seam Statistics

Features	Classification Accuracy for Benign Seam-Carving of:			
	5%	10%	20%	30%
Wavelet Absolute Moments (1 <sup>st</sup> Seam Carve)	44.07	61.05	73.65	84.30
Wavelet Absolute Moments (2 <sup>nd</sup> Seam Carve)	44.54	59.64	70.89	81.54
Wavelet Absolute Moments (1 <sup>st</sup> and 2 <sup>nd</sup> Seam Carve)	43.13	66.71	79.04	87.87

The performance of the overall combined feature set is shown in Table 5.

Table 5: Classification Accuracy of Complete Set of Combined Features Using Benign Seam-Carving

Reduction Amount	Accuracy (%)	False Positive (%)	False Negative (%)
5%	63.27	16.44	20.28
10%	73.25	13.68	13.07
20%	84.03	7.35	8.63
30%	<b>91.31</b>	4.45	4.25

The ROC curve for the overall combined feature set is shown in Figure 9. It can be seen that, while performance is clearly dependent upon the percentage of reduction, for a seam-carving of 30 percent a detection rate of 90 percent can be obtained with a false positive rate under 10 percent.




Figure 9: Receiver Operating Characteristic curves for our feature set.

### 3.3 Detection of Benign Enlargement

Benign enlargement using seam-carving identifies seam paths with minimum cumulative energy and inserts a seam along these seam paths using the average value of the seam on either side of the inserted seam<sup>10</sup>. This introduces very specific structure in the output image. In the absence of compression or other image processing, the detection of image enlargement via seam-carving is significantly easier than detection of image reduction. If the inserted seams are merely

the average of pixel values on either side of the seam, a brute force method of finding a connected path of pixels whose values correlate to their adjacent pixels in this manner can be developed. This would allow a statement regarding seam-carving to be made with very high probability. If no such connected path is found, then we can say with certainty that the image was not enlarged by seam-carving.

With this as a backdrop, we consider detection of image enlargement as a simpler problem and have concentrated our work on benign and deliberate reduction. We have tested detection performance using the set of features developed for reduction. Our results are given in Table 6.

Table 6: Classification Accuracy of Combined Features in Detection of Benign Enlargement

Enlargement Amount	Accuracy (%)	False Positive (%)	False Negative (%)
5%	92.79	3.64	3.57
10%	93.67	3.37	2.97
20%	91.85	4.04	4.11
30%	<b>95.75</b>	2.36	1.89

### 3.4 Detection of Deliberate Reduction

Detection of deliberate reduction is a rich problem which merits significant further study. The problem is somewhat difficult to frame due to the many factors involved – object size, object location, directionality of reduction (horizontal, vertical, bi-directional), and so on. An example of deliberate reduction is shown in Figure 10.



Figure 10: Image before (left) and after (right) deliberate reduction by 20%.

We estimated that a minimum of about 50 images was required in order to train the classifier. Our results are shown in Table 7. Using the same conditions (removal of second seam carve features, 20% reduction, Wavelet absolute moments), benign removal performance is 82.41%. Our performance is poorer, but this is somewhat expected as the forced nature of deliberate seam removal will have a different effect on image statistics. The forced nature of deliberate removal means that seams are no longer removed solely based upon their energy. This distorts both our energy bias and seam behavior metrics. We are also using a smaller training data set, but this is necessary due to the difference between benign and deliberate reduction on the images. If we use the benign removal image set to train our classifier, the classification accuracy for deliberately resized images decreases to 64%.

Table 7: Classification Accuracy Using Specific Features for a 20% Reduction by Deliberate Seam-Carving

Features	Accuracy (%)	False Positive (%)	False Negative (%)
Energy Bias Features	46.00	42.00	12.00
Seam Behavior Features	42.00	32.00	26.00
Wavelet Absolute Moments	68.00	6.00	26.00
Combined Features	76.00	4.00	20.00

Figure 11 provides a comparison of benign and deliberate reduction with plots similar to Figures 4 and 5. We observe that, from the standpoint of energy bias and seam behavior, deliberate reduction appears to be more similar to the original image than to an image seam-carved with benign reduction by an equivalent amount. The energy bias and seam behavior features are motivated by image statistics that are more specific to benign resizing. It is understandable that performance will be worse for deliberate reduction.

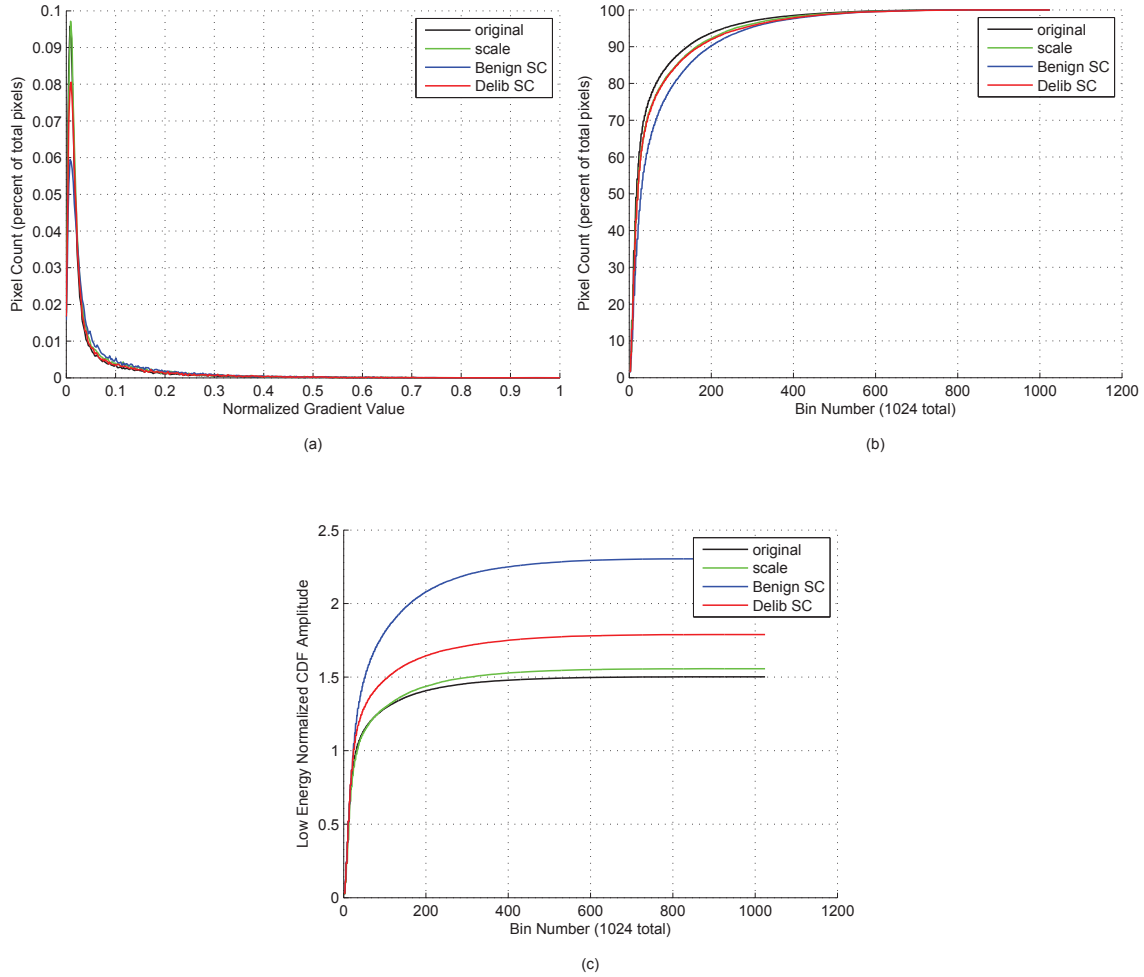


Figure 11: Comparison of energy bias features for benign and deliberate reduction by 20%. (a) 1024 bin histogram of gradient values; (b) cumulative distribution of histogram values; (c) cumulative distribution normalized by  $e_{low}$ .

## 4. CONCLUSION AND DISCUSSION

In this paper we address the problem of detecting content-adaptive scaling of images using seam-carving. We propose a set of features for the detection of seam-carving and evaluate performance of this feature set over an extensive database of images. Our results indicate that the set of features we have selected allow fairly reliable detection of seam-carving when the seam-carving operation is directed toward benign reduction or enlargement in size. Deliberate object removal via seam-carving, on the other hand, seems to confound some of the features and our classifier demonstrates only modest accuracy. Detection of deliberate reduction is a rich problem which merits significantly further study. The problem is somewhat difficult to frame due to the many factors involved including object size, object location, and directionality of reduction.

## REFERENCES

- [1] GIMP - The GNU Image Manipulation Program. Available: <http://www.gimp.org/>
- [2] Adobe Photoshop. Available: <http://www.adobe.com/products/photoshop/photoshop/>
- [3] Celik, M. U., Sharma, G., and Tekalp, A. M., "Lossless watermarking for image authentication: a new framework and an implementation," IEEE Transactions on Image Processing, 15(4), 1042-1049 (2006).
- [4] Podilchuk, C. I. and Delp, E. J., "Digital watermarking: algorithms and applications," IEEE Signal Processing Magazine, 18(4), 33-46 (2001).
- [5] Farid, H., "A Survey of image forgery detection," IEEE Signal Processing Magazine, 26(2), 16-25 (2009).
- [6] Johnson, M. K. and Farid, H., "Exposing digital forgeries in complex lighting environments," IEEE Transactions on Information Forensics and Security, 2(3), 450-461 (2007).
- [7] Popescu, A. C. and Farid, H., "Exposing digital forgeries by detecting traces of resampling," IEEE Transactions on Signal Processing, 53(2), 758-767 (2005).
- [8] Luka, J., Fridrich, J., and Goljan, M., "Detecting digital image forgeries using sensor pattern noise," Proc. SPIE 6072, 362-372 (2006).
- [9] Fridrich, J., Soukal, D., and Lukas, J., "Detection of copy-move forgery in digital images," Proceedings of DFRWS, (2003).
- [10] Avidan, S. and Shamir, A., "Seam carving for content-aware image resizing," ACM Transactions on Graphics, 26(3), 10 (2007).
- [11] Ng, T. and Chang, S., "A model for image splicing," IEEE International Conference on Image Processing, 24-27 (2004).
- [12] Liquid Rescale GIMP plugin. Available: <http://liquidrescale.wikidot.com/>
- [13] Liu, F. and Gleicher, M., "Automatic Image Retargeting with Fisheye-View Warping," ACM Symposium on User Interface Software and Technology, 153-162 (2005).
- [14] Setlur, V., Takagi, S., Raskar, R., Gleicher, M., and Gooch, B., "Automatic image retargeting," ACM Conference on Mobile and Ubiquitous Multimedia, 59 - 68 (2005).
- [15] Schaefer, G. and Stich, M., "UCID: an uncompressed color image database," Proc. SPIE 5307, 472-480 (2003).
- [16] Cortes, C. and Vapnik, V., "Support-vector networks," Machine learning, 20(3), 273-297 (1995).
- [17] Altun, O., Sharma, G., Celik, M., Sterling, M., Titlebaum, E., and Bocko, M., "Morphological steganalysis of audio signals and the principle of diminishing marginal distortions," IEEE Intl. Conf. Acoustics Speech and Sig. Proc., 21-24 (2005).
- [18] Goljan, M., Fridrich, J., and Holtyak, T., "New Blind Steganalysis and its Implications," Proc. SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VIII, 1-13 (2006).
- [19] Farid, H. and Lyu, S., "Higher-order wavelet statistics and their application to digital forensics," IEEE Workshop on Statistical Analysis in Computer Vision, 1-8 (2003).
- [20] Holtyak, T., Fridrich, J., and Voloshynovskiy, S., "Blind Statistical Steganalysis of Additive Steganography Using Wavelet Higher Order Statistics," Proc. of the 9th IFIP TC-6 TC-11 Conference on Communications and Multimedia Security, (2005).
- [21] Chang, C. C. and Lin, C. J., LIBSVM -- A Library for Support Vector Machines. Available: <http://www.csie.ntu.edu.tw/~cjlin/libsvm/>
- [22] Hsu, C. W., Chang, C. C., and Lin, C. J., A practical guide to support vector classification. July, 2003. Available: <http://ntu.csie.org/~cjlin/papers/guide/guide.pdf>