Exposing Image Splicing with Inconsistent Local Noise Variances

CS 663: Course Project

Adit Akarsh (19D070003)
Ishan Kapnadak (190070028)
Thomas Jacob (190070068)

Contributions

1. Adit Akarsh (19D070003) - Block average using integral image (dynamic programming) and vectorization, Presentation Classification
2. Ishan Kapnadak (190070028) - DCT bases computation, Presentation Results, Conclusions, and Improvements
3. Thomas Jacob (190070068) - Local noise variation estimation script, Presentation Theoretical Basis

Table of contents

Introduction

Image Splicing

- Our project considers the image tampering operation of image splicing, where a selected region of one image is inserted into another image.
- Our aim is to identify such images, and pinpoint the approximate location of the spliced region in images.

Theoretical Basis

Kurtosis concentration and Local noise variance

- The method estimates the local noise variance in different regions of the image and uses disparities from natural image trends to find regions of potential splicing.
- The estimation is based on the property of natural images to have a Kurtosis value quite close to a certain range in all band pass filtered domains like DCT.

$$\mathrm{Kurt}(X) = \frac{\mathbb{E}[(X-\mu)^4]}{(\mathbb{E}[(X-\mu)^2])^2}$$

- We utilize the concept of kurtosis concentration in approximating the kurtosis values across the different band pass filtered domains to a constant
- Minimization of the average sum of squared differences between the statistically estimated and observed kurtosis values is used to determine the kurtosis and variance of the underlying noise
- The closed forms of the optimal kurtosis and variance are

$$\sqrt{\kappa} = \frac{\left\langle\sqrt{\tilde{\kappa}_k}\right\rangle_k \left\langle \frac{1}{(\tilde{\sigma}_k{}^2)^2} \right\rangle_k - \left\langle \frac{\sqrt{\tilde{\kappa}_k}}{(\tilde{\sigma}_k{}^2)^2} \right\rangle_k \left\langle \frac{1}{\tilde{\sigma}_k{}^2} \right\rangle_k}{\left\langle \frac{1}{(\tilde{\sigma}_k{}^2)^2} \right\rangle_k - \left\langle \frac{1}{\tilde{\sigma}_k{}^2} \right\rangle_k^2}$$

$$\sigma^2 = \frac{1}{\left\langle \frac{1}{\tilde{\sigma}_k{}^2} \right\rangle_k} - \frac{1}{\sqrt{\kappa}} \frac{\left\langle\sqrt{\tilde{\kappa}_k}\right\rangle_k}{\left\langle \frac{1}{\tilde{\sigma}_k{}^2} \right\rangle_k}$$

Local noise variance estimate

- This is a simple extension of the global variance estimation concept to the pixel level
- The pixel-wise variance and the kurtosis values of the noisy image in each bandpass filtered domain is given by

$$\sigma^2 = \mu_2 - \mu_1^2$$

$$\sqrt{\kappa} = \frac{\mu_4 - 4\mu_3\mu_1 + 6\mu_2\mu_1^2 - 3\mu_1^4}{\mu_2^2 - 2\mu_2\mu_1^2 + \mu_1^4} - 3$$

- The moments for each local window are estimated using integral image, an efficient dynamic programming technique to compute sums of rectangular regions
- This is finally followed by the pixel-wise application of the $\sqrt{\kappa}$ and σ^2 formulas described for global noise variance

Classification

Classification of Spliced Images

- Our method generates the approximate pixelwise estimate of local noise variance in the image.
- Using the values of these local variances, we tried to train for an automatic classification into spliced/not-spliced images.

The Method

1. We compute the ratio of the mean of highest 10 local noise variances and median of local noise variances.
2. Using hyperparameter tuning, we find a threshold for this ratio such that the images with the ratio values below this threshold are not spliced, and the ratio values above this are spliced.

Results

Results

- We run our algorithm over the entire dataset (which has both spliced and authentic images) and classify images as being spliced or authentic.
- We then generate the confusion matrix for our classification and compute relevant classification metric.
- The confusion matrix and these computed metrics are displayed on the next slide.

Results

	Predicted No	Predicted Yes
Actual No	87	38
Actual Yes	37	83

Table 1: The Confusion Matrix

$$
\begin{aligned}
\text{Accuracy} &= \frac{170}{245} = 69.38\% \\
\text{Precision} &= \frac{83}{121} = 68.59\% \\
\text{Recall} &= \frac{83}{120} = 69.17\% \\
F_1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} &= 68.88\%
\end{aligned}
$$

(a) An authentic image

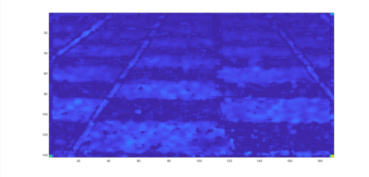

(b) its local noise variance estimate

Figure 1: We see that the local noise variance estimate does not have extremely high values and thus, the image is correctly classified as authentic.

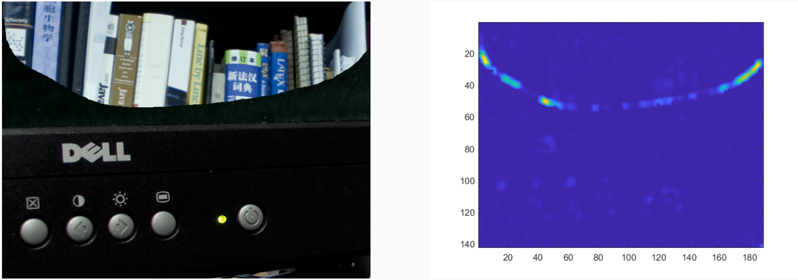

(a) A spliced image

(b) its local noise variance estimate

Figure 2: We see that the local noise variance estimate has extremely high values along the edge of the spliced region and thus, the image is correctly classified as being spliced.

Some Examples

(a) An authentic image

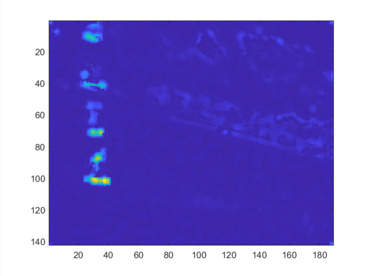

(b) its local noise variance estimate

Figure 3: Although the above image is authentic, we see that the presence of files on an ambient background disturbs the noise characteristics, leading to high noise variance values. Thus, this image was misclassified as being spliced.

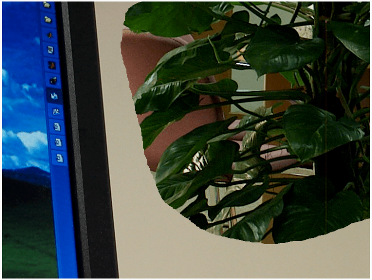

(a) A spliced image

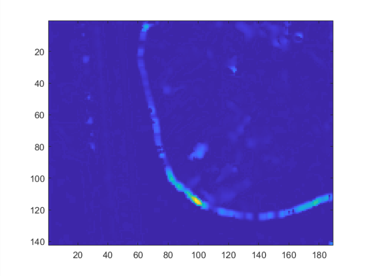

(b) its local noise variance estimate

Figure 4: We see that the local noise variance estimate has extremely high values along the edge of the spliced region and thus, the image is correctly classified as being spliced.

Some Examples

(a) An authentic image

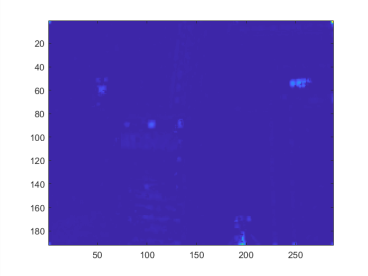

(b) its local noise variance estimate

Figure 5: We see that the local noise variance estimate does not have extremely high values and thus, the image is correctly classified as authentic.

Some Examples

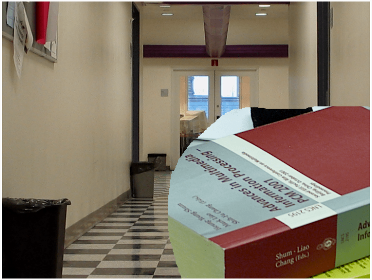

(a) A spliced image

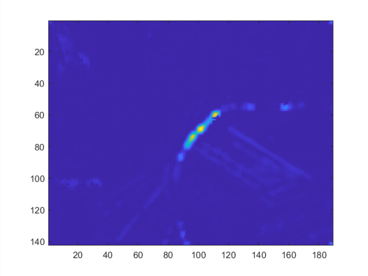

(b) its local noise variance estimate

Figure 6: We see that the local noise variance estimate has extremely high values along the edge of the spliced region and thus, the image is correctly classified as being spliced. Note that although the high noise region is faint, our method is robust enough to detect this since we only consider the top 10 values.

(a) An authentic image

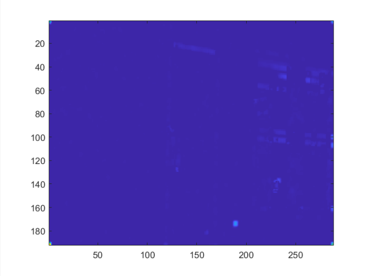

(b) its local noise variance estimate

Figure 7: We see that the local noise variance estimate does not have extremely high values and thus, the image is correctly classified as authentic.

Some Examples

(a) A spliced image

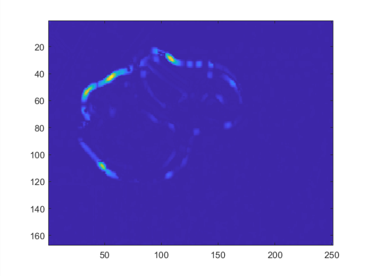

(b) its local noise variance estimate

Figure 8: We see that the local noise variance estimate has extremely high values along the edge of the spliced region and thus, the image is correctly classified as being spliced.

(a) An authentic image

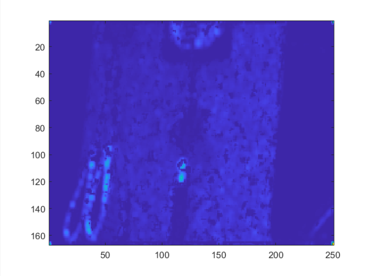

(b) its local noise variance estimate

Figure 9: We see that the local noise variance does not have extremely high values and thus, the image is correctly classified as authentic

(a) A spliced image

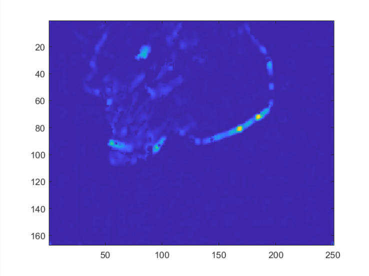

(b) its local noise variance estimate

Figure 10: We see that the local noise variance does not have extremely high values and thus, the image is correctly classified as authentic

Conclusions and Improvements

Conclusions

- We see that our method performs fairly well in classifying images as spliced or authentic, and is robust to extremely faint/thin regions of high local noise variances.
- Even in authentic images, there are some high noise variance values along object boundaries, which may 'fool' the classifier into thinking that the image is spliced. This was extremely predominant in the Windows Home Screen example. We were not able to find a suitable threshold value to accurately discriminate between spliced regions and object boundaries.

Future Improvements

- Our method of using a threshold for the ratio is quite susceptible to object boundaries, and thus is not very accurate. An improvement we had in mind was to use an edge detection algorithm such as Canny Edge Detection on the noise variance image. If we are able to accurately detect edges in the noise variance, we should also be able to discriminate between spliced regions and object boundaries.
- In our implementation, we have only used a DCT as the pre-smoothing filter. The paper also mentions the wavelet decomposition as an alternate pre-smoothing filter. As part of our future work, we may experiment with different types of pre-smoothing filters and see which gives us the best results.