

FACE LIVENESS DETECTION UNDER BAD ILLUMINATION CONDITIONS

Bruno Peixoto, Carolina Michelassi, and Anderson Rocha

University of Campinas (Unicamp)
Campinas, SP, Brazil

ABSTRACT

Spoofing face recognition systems with photos or videos of someone else is not difficult. Sometimes, all one needs is to display a picture on a laptop monitor or a printed photograph to the biometric system. In order to detect this kind of spoofs, in this paper we present a solution that works either with printed or LCD displayed photographs, even under bad illumination conditions without extra-devices or user involvement. Tests conducted on large databases show good improvements of classification accuracy as well as true positive and false positive rates compared to the state-of-the-art.

Index Terms— Liveness Detection, Face Spoofing Detection, Counter-Spoofing

1. INTRODUCTION

Biometrics refers to technologies designed either to verify or recognize the identity of a person based on one or more physical or behavioral characteristics, and has been widely used in authentication systems. Such methods, however, can be fooled (spoofed) by an identity thief, specially the ones based on face recognition, where the thief can obtain a photo of an authentic user from a significant distance, or even obtain it from the Internet.

In an operational scenario, a 2-D image-based facial recognition system can be spoofed by simple tricks. For instance, instead of showing one's own face to the biometric system, an unauthorized person can display a photo of an authorized counterpart either printed on a piece paper, on a laptop, or even on a cell phone screen.

Most of the current face recognition research efforts strives for dealing with the "image matching" part of the system. Consequently, some poorly-designed systems have been shown to be fooled by very crude line drawings of a human face [1]. Although there has been important advances regarding spoofing detection, this research branch is still an open problem. In addition, non-intrusive methods without extra devices and human involvement are preferable in practice, given that they could be integrated into an existing face recognition system, where usually only a generic webcam is deployed [2].

In this paper, we extend a very powerful spoofing detection system recently introduced by Tan et al. [2]. In their work, the authors present a face liveness detection system using a single image. Although the authors presented effective results, in this paper we extend their work so as to be robust to bad illumination conditions. Our experiments performed on the same data set as Tan et al's present more than 50% reduction in the classification/detection error.

We thank the São Paulo Research Agency (FAPESP) for funding the research under the Award 2010/05647-4. We also express our gratitude to X. Tan and colleagues for kindly providing us with the data set and promptly answering our e-mails.

Corresponding Author: anderson.rocha@ic.unicamp.br

2. STATE-OF-THE-ART

Some anti-spoofing techniques for facial recognition systems rely on Fourier analysis, which can differentiate between live faces and certain types of spoofs, such as printed images, based on the differences in the high frequencies of the Fourier spectra [2, 3].

Another branch of research has focused on motion detection such as eye blinking and involuntary movements of parts of the face and head [4, 5]. However, these spoofing detection methods are still impacted by simple head tilts simulating head movement or by short video sequences displaying an authentic user. Other used facial cues include the surface texture of the skin. With such an analysis, we can calculate certain measures to characterize the optical qualities of the facial skin of live people and compare to non-live ones [6].

If extra devices are available, near infrared or thermal images can be considered [7]. The 3-D information can also be used to provide additional protection against spoof attempts over such devices such as 3-D cameras or multiple 2-D cameras [8]. If there is user involvement, we can take advantage of multi-modal information (e.g., voices or gesture) and various challenge-response methods (e.g., asking the user to blink, or smile) [2].

Besides the face liveness detection, one can focus on the security of biometric templates using a user-specific key to generate a random matrix to distort the face template, so that a "stolen" face image without the key will be almost of no use [9].

Although there has been important advances regarding spoofing detection such as the aforementioned ones, this research branch is still an open problem. Two of the most important challenges nowadays refer to: (1) the need of designing and deploying non-intrusive methods without extra devices and human involvement; and (2) designing detection methods robust to changes in pose and illumination. In this paper, we deal with the latter problem: to detect spoofs under challenging illumination conditions without the need of extra devices or user involvement.

3. PROPOSED SOLUTION

One way to improve liveness detection without extra devices or user involvement is to recognize when the image (photo or video) comes from a laptop or cell phone screen instead of a real human [2].

Although Tan et al. have presented very effective results in their most recent work [2], the authors overlooked a problem often found in an operational scenario: bad illumination conditions.

In this paper, we extend their method to deal with images even under hard illumination conditions either for spoof attempts coming from a laptop display or high-quality printed images.

3.1. The Basic Model

The image of a photograph shown on an LCD screen taken from a camera is intuitively the image of a real face, but passes through the

camera system twice and the projection display once. The brightness of the LCD screen affects the recaptured image in a way that the high-frequency regions (borders) become susceptible to a “blurring” effect due to the pixels with higher values brightening their neighborhood. This makes the recaptured images show less borders than their real counterparts.

In this paper, we explore such information in order to detect whether an image is a spoof or not. For that, we analyze the image similar to [2], using a Difference of Gaussian (DoG) filter, which is a bandpass filter that uses two Gaussian filters with different standard deviations as limits. The idea is to keep the high-middle-frequencies to detect the borders, but not much, in order to remove the noise (which is also a high-frequency). Figure 1 depicts an image that passes through this filter. In this work, we used the same standard deviations, $\sigma_1 = 0.5$ and $\sigma_2 = 1.0$ proposed by Tan et al. [2].



Fig. 1. A DoG filtering example. The original image (far left), its DoG, the recaptured image (third column) and its DoG (far right). Notice that the first DoG has more details than the second one.

For the classification stage, we used the standard *Sparse Logistic Regression Model* similar to the first model in Tan et al. [2]. This model directly uses the images filtered with the Difference of Gaussian (DoG) method and transformed to the frequency domain as feature vectors.

The Sparse Logistic Regression Model works as follows: let $\mathbf{x} \in \mathbb{R}^n$ denote a sample, which is the linearized version of the image under investigation (with $n = w \times h$ pixels), and $y \in \{-1, 1\}$ the associated label defined as ‘+1’ for the recaptured image (i.e., spoof) and ‘-1’ for the real image (i.e., non-spoof or live image). The logistic regression model is given by

$$\text{Prob}(y|\mathbf{x}) = \frac{1}{1 + \exp(-y(\mathbf{w}^T \mathbf{x} + b))}, \quad (1)$$

where \mathbf{w} and b are the values obtained from the logistic loss minimization (training step) and $\text{Prob}(y|\mathbf{x})$ is the conditional probability of class $y = +1$ given a sample \mathbf{x} , meaning the sample \mathbf{x} is a recaptured image. The average logistic loss is given by

$$\text{loss}(\mathbf{w}, b) = \frac{1}{m} \sum_{i=1}^m \log(1 + \exp(-y_i(\mathbf{w}^T \mathbf{x}_i + b))), \quad (2)$$

which is a smooth and convex function. We can determine \mathbf{w} and b by minimizing the average logistic loss: $\min_{\mathbf{w}, b} \text{loss}(\mathbf{w}, b)$, leading to a smooth convex optimization problem. There are efficient solvers for optimizing this problem. In our case, we solve such problem for w and b using the logistic loss method with the L1-norm regularization in SLEP Package [10]. Please refer to [10] for more details about such an optimization problem.

The above model yields good classification results. Notwithstanding, it is not well suited to directly analyze images under bad illumination conditions given that the DoG filtering is not enough to rule out the illumination artifacts and properly detect the image borders. We have found that different parts of an image behave differently under illumination variations (e.g., suffering partial occlusions and shadows) and a standard histogram normalization is not

enough to solve such problems. Therefore, we need to analyze image regions adaptively correcting each region according to their local properties in order to maximize the classification detection rate at the final stage.

3.2. Dealing with Bad Illumination

In an operational scenario, the illumination conditions vary greatly, many times causing shadows on parts of the face. These darker regions of the image may influence the spoof detection method we described in Section 3.1, mainly because the DoG filtering does not detect the borders properly under bad illumination conditions (See Fig. 2 for an example).




Fig. 2. Example of an image after the DoG filtering stage for the same person under good illumination (left) and bad illumination conditions (right). The bad illumination causes a shadow on the left half of the face preventing DoG filtering from finding the correct edges.

To minimize this effect and improve the feature discriminability, we pre-process the image in order to homogenize it as a whole, so their borders could be easier detected and the illumination changes become more controlled. For that, the first obvious solution is to use a simple histogram equalization. However, this method is highly sensitive to images with unbalanced dark and bright parts, such as many images of one of the difficult data sets we use in this paper. The *Yale Face Database B* contains images whose bright parts get even brighter, losing edges, and the dark ones remain dark, keeping the image unbalanced and decreasing the edges response found by the DoG filter. Figure 3 shows an example of this effect.

Given that the illumination changes affect the image with different degrees and in a localized manner, in this paper we use the contrast-limited adaptive histogram equalization (CLAHE) [11] approach to deal with these changes. CLAHE operates on small regions in the image, called tiles. Each tile’s contrast is enhanced, so that the histogram of the output region approximately matches a specified distribution: **Uniform** (resulting on a flat histogram), **Rayleigh** (resulting on a bell-shaped histogram), or **Exponential** (resulting on a curved histogram). For instance, using the Rayleigh distribution, the updated value of a pixel p considering a tile with cumulative probability distribution $P(f)$ is given by $p_{new} = (2\alpha^2 \ln(\frac{1}{1-P(f)}))^{-0.5}$.

The neighboring tiles are then combined using bilinear interpolation to eliminate artificially induced boundaries. The contrast can be limited to avoid amplifying eventual noise and to prevent over-saturation of the image, specifically in homogeneous areas. Without this limit, the technique could produce results that, in some cases, are worse than the original image. Figure 4 shows an example of an image that goes through this adaptive histogram equalization.

The method relies on the choice of a three parameters: (1) number of tiles (higher values lead to better results at the cost of a longer computational effort); (2) contrast enhancement limit (higher values result in more contrast); and (3) alpha value, α (related to the



Fig. 3. Simple histogram equalization (middle) and the respective DoG filter response (far right) of the input image on the left.




Fig. 4. Image with adaptive histogram equalization and its DoG.

Rayleigh and Exponential distributions). In the spoofing detection scenario, a good tradeoff on these parameters can easily be calculated using a small training set with normal and recaptured images under varying illumination changes.

4. EXPERIMENTS AND DISCUSSION

4.1. Databases

To validate the proposed ideas and compare results to the state-of-the-art, we used two data sets: the first one is the *Yale Face Database B* [12], which comprises 640 real faces, taken from 10 test subjects with 64 different illumination conditions. The use of this data set is explained below. The other data set is the *NUAA Imposter Database*, used by Tan et. al. [2], which comprises 5,105 real faces and 7,509 imposter faces (real faces recaptured after printing). We used Tan et. al's (*NUAA Imposter*) data set to validate our contribution and compare it to Tan et. al's. [2] original results.

To investigate the spoofing detection method with respect to recaptured images from LCD monitors, we set up a recapture experiment to collect images from a 50cm-distance from an LCD screen. For that, we recaptured the Yale Face Database B using three LCD monitors, an LG Flatron L196WTQ Wide 19", a CTL 171Lx 17" TFT and a DELL Inspiron 1545 notebook. The cameras used were a Kodak C813 8.2 megapixels and a Samsung Omnia i900, with 5 megapixels. After the recapture, we cropped the images to center the faces, resampled to 64×64 pixels and converted them to grayscale. The final data set comprises 640 real faces and 1,920 LCD spoofs (available under request: <http://www.ic.unicamp.br/~rocha/pub/communications.html>).

4.2. Experiment Setup

In all experiments in this paper, we have performed 10-fold cross validation. Each feature vector representing an image to be classified contains the image's linearized pixels processed according to the approaches we discussed in the previous sections.

All the training samples are used in the SLEP package [10], and minimized to yield a weight vector \mathbf{w} and the intercept b of the lo-

gistic loss classification method we use. To test an input example, we simply calculate Equations (1) and (2) using the weight vector \mathbf{w} and the intercept b learned in the training stage.

4.3. Results

As we discussed in Section 3.2, the illumination process requires the selection of the number of tiles we want to use to normalize the image locally, the contrast enhancement limit and an α value when the chosen distribution is other than the Uniform. One can easily find a good parameter tradeoff automatically using the training examples. Table 1 shows the Equalization parameters tested.

Table 1. Parameters used for equalization

Distributions	Uniform	Exponential	Rayleigh		
# Tiles	4	6	8	10	
Contrast limit	0.01	0.03	0.05	0.07	0.1
α (Alpha)	0.3	0.5	0.7	0.9	

Table 2 shows the experiment results for the *NUAA Original Database* provided by [2]. Excellent results are obtained with the proposed extension using Rayleigh distribution operating with eight tiles, contrast limit of 0.07 and $\alpha = 0.5$. Recall that Table 1 shows the range of possible values for the normalization parameters which can easily be computed using the provided training examples. This experiment shows that the proposed extension captures even slight changes in illumination (NUAA database only not contains small illumination changes). The proposed method achieves a classification accuracy mean of $\mu \cong 93\%$ or just $\epsilon \cong 7\%$ classification error while Tan et al. [2] yields a classification error of $\epsilon \cong 13.4\%$ (more than 50% in error reduction).

As a matter of fact, we also observed that using the same Rayleigh distribution with lower values of α tend to yield worse classification results. With $\alpha = 0.3$ the intermediate values of the pixels tend to be lower, darkening the image and making it more different to the original image. We also observed that using less tiles and a contrast limit superior to 0.1 are also not interesting. The equalization method with four tiles splits the image into big regions, and may have the same problems with non-adaptive equalization, which uses the entire image.

Table 2. Experiment results for NUAA Imposter Database. Tan et al [2] approach v. Proposed Extension for Bad Illumination Conditions.

Tan et al's Approach [2]				
	Min	Mean	Max	STD
Classification Accuracy	85.2%	86.6%	87.5%	0.6%
True Positive Rate	81.9%	82.4%	90.4%	0.6%
False Positive Rate	8.0%	9.3%	18.8%	1.3%
Proposed Extension — Bad Illumination Conditions				
8 tiles, Rayleigh Distribution, Contrast Limit = 0.07, $\alpha = 0.5$.				
	Min	Mean	Max	STD
Classification Accuracy	92.0%	93.2%	94.5%	0.4%
True Positive Rate	92.6%	93.0%	93.7%	0.4%
False Positive Rate	4.7%	6.7%	8.4%	1.3%

Table 3 presents the classification results for the *Yale Recaptured Database*. This is a database with very difficult illumination conditions which contains normal as well as recaptured images from LCD screens (see Section 4.1). The data set comprises darker images than in the NUA database, many of them having some parts of the face not visible due to sideways illumination effects.

Table 3. Experiment results for Yale Recaptured Database. Tan et al [2] approach v. Proposed Extension for Bad Illumination Conditions.

Tan et al's Approach [2]				
	Min	Mean	Max	STD
Accuracy	71.0%	76.1%	81.2%	1.0%
True Positive Rate	92.0%	93.6%	95.1%	1.0%
False Positive Rate	32.8%	41.4%	50.0%	7.4%

Proposed Extension — Bad Illumination Conditions				
4 tiles, Rayleigh Distribution, Contrast Limit = 0.1, $\alpha = 0.1$.				
	Min	Mean	Max	STD
Accuracy	89.0%	91.7%	93.8%	1.4%
True Positive Rate	83.3%	85.8%	87.7%	1.4%
False Positive Rate	0.00%	2.5%	5.4%	1.5%

With the original method (without dealing with illumination changes), the classifier achieves a mean classification accuracy of $\mu \cong 76\%$ with a high false positive rate ($FPR \cong 41\%$) which is unacceptable for this type of application that should not deny access for authorized users. When we cope with the bad illumination condition, the classifier yields a surprising mean classification accuracy of $\mu \cong 91.7\%$ with $FPR \cong 2.5\%$. The proposed extension reduced the classification error in more than 65%. We obtain such results using the Rayleigh distribution which better preserves the original aspects of the image across its bell-shaped histogram modeling.

We have observed that the Exponential distribution also provides good results. However, sometimes it excessively brightens the originally well-lit regions of the image, causing the DoG filter to fail at detecting the borders of these regions.

In the *Yale Recaptured Database*, the α value that presented the best results was 0.3. Given that the images of this database are darker, a value too high of alpha leads to brighter images, distorting their original aspects. We also have observed that higher contrast limits often lead to better classification results. Since we are trying to detect the high-frequency regions, a higher contrast of the image helps the filter to detect them, especially on images that present regions too dark.

The number of tiles dictates how the image is divided for the adaptive equalization. Since many images in the *Yale Recaptured Database* have approximately half of the face lit and half shadowed, each distribution reacts differently to the number of tiles. However, for all these cases, a few training examples are enough to learn suitable parameter values.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we dealt with the problem of spoofing detection in face recognition systems using images under uncontrolled illumination conditions. We also analyzed spoofing attempts using high quality printed images (the problem tends to be simpler with low-quality

print attacks) as well as recaptured images from LCD screens. Although, there are other possible forms of spoof (e.g., video-based), here we dealt with recognition systems based on only one image.

The proposed extension method to detect spoofing attempts under illumination changes achieves good classification results with low false positive and false negative rates even with dark images. The proposed extension reduced the classification error in more than 50% for high-quality printing spoofs (*NUAA Imposter Database*) and more than 65% for the case of recaptured LCD images (*Yale Recaptured Database*) when compared to the state-of-the-art counterpart.

As a future work, we seek to improve the parameter search during training. We also are interested in testing other face-based spoofs such as the ones coming from cell-phone screens.

Finally, currently we are pursuing some approaches for dimensionality reduction in order to take advantage of powerful classification methods such as Support Vector Machines (SVMs).

6. REFERENCES

- [1] K. Nixon, V. Aimale, and R. Rowe, "Spoof detection schemes," in *Handbook of Biometrics*. Springer US, 2008.
- [2] X. Tan, Y. Li, J. Liu, and L. Jiang, "Face liveness detection from a single image with sparse low rank bilinear discriminative model," in *ECCV*, 2010.
- [3] J. Li, Y. Wang, T. Tan, and A. K. Jain, "Live face detection based on the analysis of fourier spectra," in *Biometric Technology for Human Identification*, 2004, pp. 296–303.
- [4] G. Pan, Z. Wu, and L. Sun, *Recent Advances in Face Recognition*, chapter Liveness detection for face recognition, pp. 235–252, InTech, 2008.
- [5] K. Kollreider, H. Fronthaler, and J. Bigun, "Non-intrusive liveness detection by face images," *Image Vision Comput.*, pp. 233–244, February 2009.
- [6] G. Parziale, J. Dittmann, and M. Tistarelli, "Analysis and evaluation of alternatives and advanced solutions for system elements," *BioSecure*, 2005.
- [7] D. A. Socolinsky, A. Selinger, and J. D. Neuheisel, "Face recognition with visible and thermal infrared imagery," *CVIU*, pp. 72–114, July 2003.
- [8] T. Fladsrud, "Face recognition in a border control environment: Non-zero effort attacks' effect on false acceptance rate," Master thesis, Gjøvik University College, Norway, 2005.
- [9] A. Teoh, A. Goh, and A. Ngo, "Random multispace quantization as an analytic mechanism for bihashing of biometric and random identity inputs," *IEEE TPAMI*, pp. 1892–1901, 2006.
- [10] J. Liu, S. Ji, and J. Ye, "SLEP: Sparse learning with efficient projections," 2009, Available at <http://www.public.asu.edu/~jye02/Software/SLEP>.
- [11] K. Zuiderveld, "Contrast limited adaptive histogram equalization," *Graphic Gems IV*, pp. 474–485, 1994.
- [12] A. S. Georghiadis, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE TPAMI*, pp. 643–660, 2001.