

Automatic Smile and Blur Detection for Still Images

Yossi Amoyal and Evgeny Gershikov
Department of Electrical Engineering
Ort Braude Academic College of Engineering
Karmiel 21982, Israel

Abstract—In this work we present an algorithm designed to recognize if the faces in a still image are smiling or not, based on Viola Jones face detection method and an algorithm for smile recognition. The stages of our algorithm are: identification of the faces in the image, detection of the mouth part of each face and then application of a basic algorithm to the mouth part to determine whether the person is smiling or not. A smile is recognized by examining the graph of the grayscale levels across the mouth and looking for two distinct minima. As an additional step of image quality assessment, we measure the level of blurriness of the detected face. Real life images were used for testing the algorithms at different spatial resolutions, levels of blurriness as well as varying contrast and brightness conditions. Our conclusion is that the proposed methods can be beneficial for automatic smile recognition and blur detection in digital cameras or other mobile devices.

Keywords—smile recognition, blur detection, face image processing, automatic image quality assessment

I. Introduction

Today many consumer electronic devices can be used to capture digital pictures, for example, digital cameras, mobile computers or mobile phones. For such devices, smile detection may be employed in various applications. One possible application is that smile detection can be used in digital cameras or camera phones to automatically determine when the shutter should be closed or when a taken picture should be discarded. Another possible application of smile detection is to conduct automatic surveys using digital signage displays. For instance, smile detection may be used to determine how many people enjoy particular movie trailers, television programs or advertisements.

In embedded applications, the computing and memory resources are usually very limited. Moreover, smile detection software may be required to be resident in memory for continuous execution (such as in digital signage). This makes power consumption an important factor. Unfortunately, existing smile detection approaches usually require significant computing and memory resources.

A number of smile detection methods are known, based on Sobel filtering, edge detection and segmentation [6]. In this work we demonstrate the use of Viola-Jones algorithm [13] for mouth area detection and implement a smile detection

algorithm applied to these area. The algorithm is based on mathematical modeling of the graph of grayscale levels along a vertical profile of the mouth [12]. It is described next.

II. Smile detection algorithm

The used method for detecting whether a person is smiling or not is based on the assumption that when a person smiles, the teeth are shown. In the next subsections we describe the used detection technique.

A. Using regions of interest

In order to speed up the algorithm and improve its accuracy as well, we propose to work in small regions of interest (ROIs) of the image that contain the mouth areas. The mouth areas can be located using the Viola-Jones method [13]. Taking one step further we capture the center of the mouth region, which for most tested images is the area of 20% of the width and 60% of the height of the mouth part around its center point.

B. Analysing of the vertical mouth profile

We use the following stages to calculate the vertical mouth profile and determine if it fits a smiling mouth profile or not. Whether smiling or not we assume the mouth is closed.

1. Calculate the average of each row of the small ROI in the mouth area. This is the gross vertical profile of the mouth.
2. Use polynomial approximation to model it as function of the vertical spatial coordinate.
3. Get the values for the maximum and the minimum or minima of the profile.
4. If there are two minima, compare the difference between their values relative to the full scale of the profile (the difference between maximal and minimal values) to a predefined threshold (e.g., 35%). If the relative difference is below the threshold, the algorithm decides that a smile is detected. In the opposite case or when only one distinct minimum is present, the algorithm's outcome is that there is no smile.

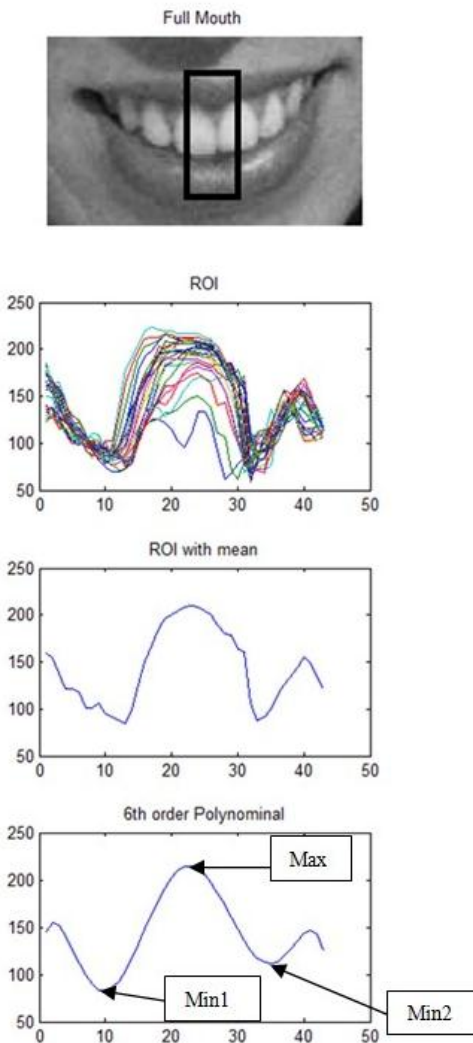


Figure 1: Analysing the profile of a smiling mouth.

Note that the algorithm stages were carefully designed to render a noise tolerant and robust algorithm. Hence the averaging of the rows in Stage 1, the polynomial approximation in Stage 2 and the thresholding operation in Stage 4, that is used to determine if the profile has indeed two distinct minima or only one.

Detection example – smiling mouth profile

An example of a smiling mouth and the profile resulting from its processing by our algorithm is given in Fig. 1. The profiles of individual columns are seen in the 2nd figure and the averaged profile in the 3rd figure from the top. The bottom figure shows the mouth profile after polynomial approximation. We can see that there are two minima in the profile and both have similar values. Thus, by comparing the difference in their values to the full range (minimum to maximum), we can detect a smile. The two minima that are seen in the profile correspond to the two dark regions that are between the upper lip and the teeth and the lower lip and the teeth.

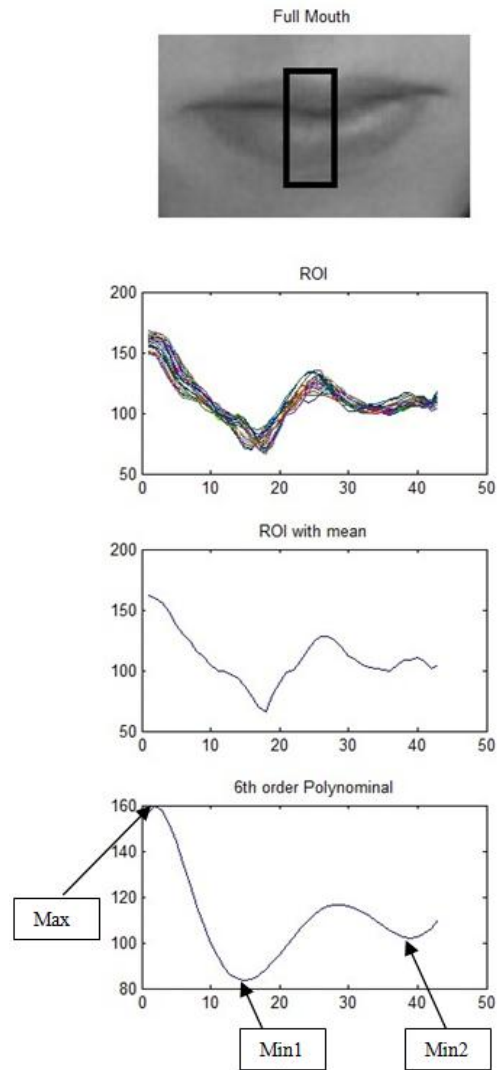


Figure 2: Analysing the profile of a non smiling mouth.

Detection example - non smiling mouth profile

Another example of a close mouth and its profile is given in Fig. 2. We can see that in this case there is one minimum that has significantly lower value than the other, corresponding to the dark region between the lips. After calculating the value of the relative difference of Step 4 in the smile detection algorithm (Section II.B.) we conclude that it is above the set threshold and thus the decision is that there is no smile.

III. Blur detection algorithm

The used method to detect the level of blurriness in the image is based on edge detection [1]. There are many kinds of edge detection filters such as Sobel [5], Gabor [11], Prewitt [10], Laplacian of Gaussian [8], Canny filter [2] and another method based on generating the edge image from the Haar Wavelet Transform [3]. Next we describe

the Laplacian of Gaussian method briefly and then in Section III.B. we present our method.

A. Laplacian of Gaussian

The Laplacian $L(x,y)$ of an image with (continuous) pixel intensity values $I(x,y)$ is given by:

$$L(x,y) = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2} \quad (1)$$

For a discrete image, this can be approximated using a convolution filter [4] $h(k,l)$:

$$L(i,j) = \sum_{k=1}^m \sum_{l=1}^n I(i+k-1, j+l-1)h(k,l), \quad (2)$$

where i changes in the range between 1 and $M-m+1$ and j changes between 1 and $N-n+1$.

The discrete convolution kernel $h(k,l)$, that can approximate the second derivatives [7] in the definition of the Laplacian, as commonly used, is shown in Fig. 3.

0	-1	0
-1	4	-1
0	-1	0

Figure 3: the discrete Laplacian filter.

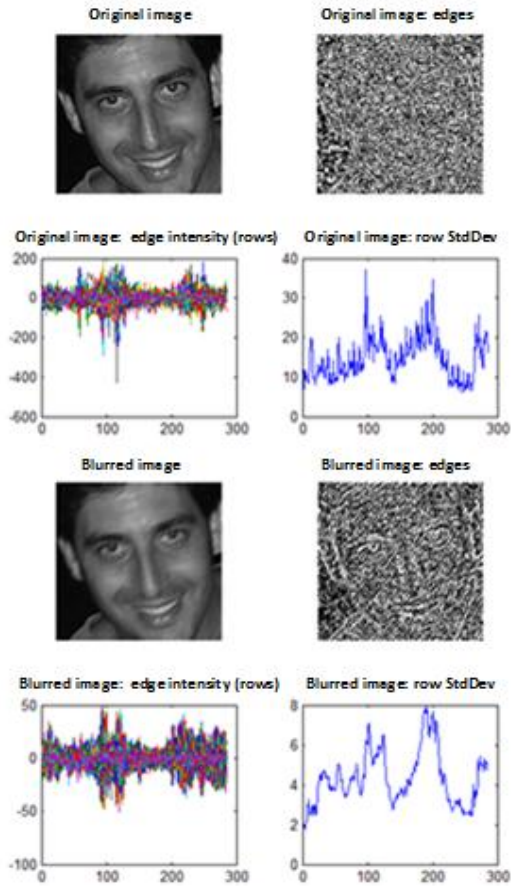


Figure 4: Analysing a face image in sharp focus: original image and its blurred version.

B. Stages of the proposed method

We use the following stages to calculate a vertical head profile and determine if it fits the profile of an out of focus or motion blurred face.

1. Apply the Gaussian filter to the image I followed by the Laplacian filter (Fig. 3).
2. Remove the border of the filtered image ($m \times n$ image will become $(m-1) \times (n-1)$ image).
3. Calculate the standard deviation (StdDev) of the rows of the filtered image.
4. Find the maximum of the StdDev and compare it to an empirical threshold. If the StdDev is above the threshold, the image is blurry.

Detection example – sharp image profile

In this example (Fig. 4) we can see the maximum of the StdDev is around 38 and its graph appears to be jagged. We can also see significant changes between the original image and its blurred version that was created with a motion blur filter, including the edge detection results.

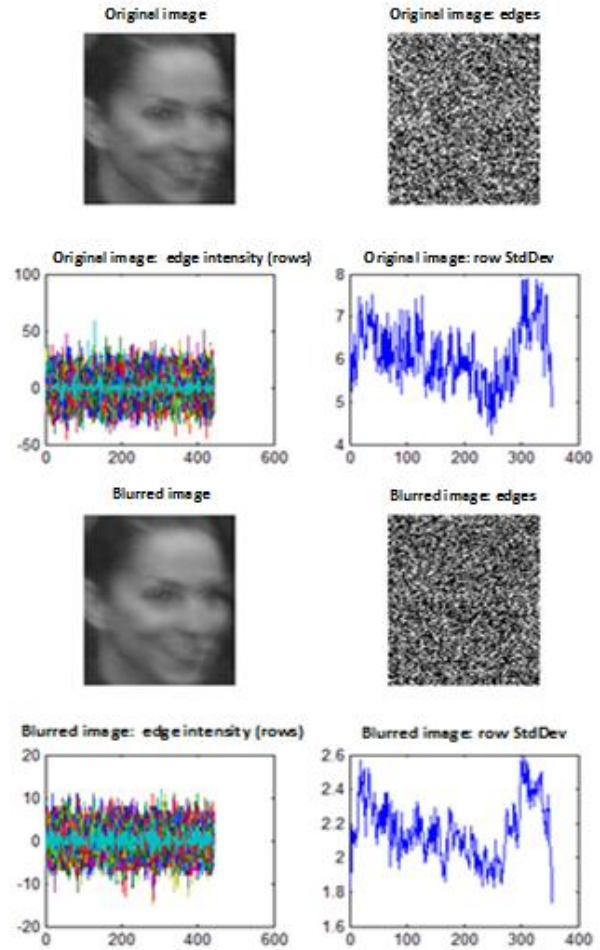


Figure 5: Analysing a blurry face image: original image and its blurred version.

Detection example – blurred image profile

A blurred image and its profile are shown in Fig. 5. Here we can see the maximum of StdDev is around 8 and its graph appears to be smoother. We can also see that there are slight changes between the original image and the blurred image (created using a motion blur filter). This is true also for the edge detection results.

IV. Testing methods and criteria

Simulations were performed, where the smile and blur detection algorithms were applied to different real life photographic images. The images taken were of high resolution since for good performance of the methods, the mouth part should be at least 180×110 pixels.

A. Images used for testing the algorithm

14 images were tested with more than 50 faces, varying brightness and contrast levels and different degrees of blurriness. Image input format was the commonly used JPEG [14]. Resolutions varied from 600×800 to 3888×2592 and most images contained two faces or more of different sizes and tilt angles. We did not examine the effect of makeup, such as different kinds of lipstick, or the effect of external objects, such as lip piercing, on the algorithm performance. We leave this for future research. Also the detection of smiles with an open mouth might be improved and is under research.

B. Testing criteria

The algorithms were tested with respect to smile and blur detection accuracy as well as speed. We measured the detection rate (DR) and false positive rate (FPR) for the smile recognition method and the total accuracy of the blur detection method. The latter algorithm chose from two levels of blurriness: low and medium/high and it is considered correct when its decision corresponds to the level of blurriness as seen visually. The algorithm speed was measured in terms of run time.

V. Results

A. Performance of the smile recognizer

The accuracy of the smile recognition [9] method in terms of DR and FPR is presented in Table 1. We can see that the algorithm accuracy is correlated with the image resolution: it increases when the image resolution increases. Run-time measurements of the algorithm in a MATLAB environment are given in Table 2. As expected, the run-time rises with the image resolution since more pixels are to be processed. We should note that the algorithm memory consumption also increases with image size and may be a significant consideration for large images.

TABLE 1: SMILE DETECTION ALGORITHM ACCURACY

Image Resolution	DR	FPR
small $\sim(800 \times 600)$	60%	15%
medium $\sim(2000 \times 1500)$	70%	10%
big $\sim(3200 \times 2400)$	78%	8%

TABLE 2: SMILE DETECTION ALGORITHM RUN-TIME

Image Resolution	Time (sec)
small $\sim(800 \times 600)$	1.5
medium $\sim(2000 \times 1500)$	3.2
big $\sim(3200 \times 2400)$	12

An additional example of correct recognition of a smiling mouth is given in Fig. 6 (left side). Another example in Fig. 6 (right side) shows the correct processing of a non-smiling mouth.

B. Performance of the blur detector

The accuracy of the blur detection algorithm was measured for a group of 72 images, some of which were in sharp focus, while others were blurry due to motion or being out of focus. The accuracy results are given in Table 3. When considering the whole group, the algorithm determined the correct level of blurriness (low or medium/high) for 97.2% of the images.

TABLE 3: BLUR DETECTION ALGORITHM ACCURACY

Test images Set	Number	Accuracy (%)
Not blurred	30	96.6
Motion blurred	26	96.2
Out of focus	16	100
Total	72	97.2

TABLE 4: BLUR DETECTION ALGORITHM RUN-TIME

Image Resolution	Time (sec)
small $\sim(800 \times 600)$	0.3
medium $\sim(2000 \times 1500)$	0.8
big $\sim(3200 \times 2400)$	1.2

VI. Conclusions

We presented two algorithms in this work, designed to recognize smiling faces and the blurriness level in still images. The faces in the tested images were detected using the Viola-Jones method and then processed by the proposed techniques. The smile recognizer consists of a stage of calculating the vertical profile of the mouth part of the face, followed by polynomial regression and detection of the number of significant minima in the profile. If there are two such minima present, the algorithm decides that the mouth is smiling, otherwise the decision is that there is no smile.

The proposed blur detector calculates the Laplacian of Gaussian (LoG) of each detected face and creates its

vertical standard deviation profile. It then compares the maximum of the standard deviation (StdDev) relative to a set threshold to determine if the image is blurred or not.

We tested the algorithms on real life photographs of people at different spatial resolutions, levels of blurriness as well as varying contrast and brightness conditions. We noticed that the performance of the smile recognizer improves significantly for larger images compared to smaller ones. The blur detector performance was excellent for almost all of the tested images. Our conclusion is that the proposed techniques should be useful for automatic face processing in applications that required smile recognition or quality assessment of a facial image.

Acknowledgment

We would like to thank the administration of Ort Braude academic college and the Department of Electrical Engineering for providing the opportunity to conduct this research and the financial means to present it at the conference.

References

- [1] M. B. Ahmad and T. S. Choi, "Local Threshold and Boolean function based Edge Detection", IEEE Transactions on Consumer Electronics, vol. 45, no. 3, pp. 674-679, August 1999.
- [2] J. Canny, "A computational approach to edge detection", IEEE Transactions on PAMI, vol. 8, no. 6, pp. 679-697, 1986.
- [3] S. A. Chatzichristofis, Y. S. Boutalis and A. Arampatzis, "Accelerating Image Retrieval Using Binary Haar Wavelet Transform on the Color and Edge Directivity Descriptor", Proc. of ICCGI, pp. 41-47, Valencia, Spain, 2010.
- [4] J. S. Chen, "Fast Convolution with Laplacian of Gaussian Masks", IEEE Trans. on PAMI, vol. PAMI-9, no. 4, pp. 584 -590, 1987.
- [5] K. Engel, M. Hadwiger, J. M. Kniss, C. Rezk-Salama and D. Weiskopf, Real-time volume graphics, AK Peters, pp. 112-114.
- [6] N. Eveno, A. Caplier and P.-Y. Coulon, "A New Colour Transformation for Lips Segmentation", IEEE Workshop on Multimedia Signal Processing (MMSP), Cannes, France, 2001.
- [7] R. M. Haralick, "Digital Step Edges from Zero Crossing of Second Directional Derivatives", IEEE Trans. on PAMI, vol. PAMI-6, no. 1, pp. 58-68, 1984.
- [8] A. Huertas and G. Medioni, "Detection of intensity changes with sub-pixel accuracy using laplacian-gaussian masks", IEEE Trans. on PAMI, vol. PAMI-8, no. 5, pp. 651-664, Sept. 1986.
- [9] A. Ito, X. Wang, M. Suzuki and S. Makino, "Smile and Laughter Recognition using Speech Processing and Face Recognition from Conversation Video", Proc. International Conference on Cyber-worlds, Singapore, 2005.

- [10] R. Maini and J. S. Sohal, "Performance Evaluation of Prewitt Edge Detector for Noisy Images", GVIP Journal, vol. 6, no. 3, pp. 39-46, 2006.
- [11] R. Mehrota, K.R. Namuduri and N. Ranganathan, "Gabor filter-based edge detection", Pattern Recognition, vol. 25, no. 12, pp. 1479-1494, December 1992.
- [12] T. Sharon, "The Smile Detector", IBM Almaden Research Center, San Jose, California, 2001.
- [13] P. Viola and M. J. Jones, "Robust Real-Time Face Detection", International Journal of Computer Vision, vol. 57, no. 2, pp. 137-154, May 2004.
- [14] G. K. Wallace, "The JPEG still picture compression standard", IEEE Trans. Consumer Electronics, vol. 38, no. 1, pp. xviii-xxxiv, 1992.

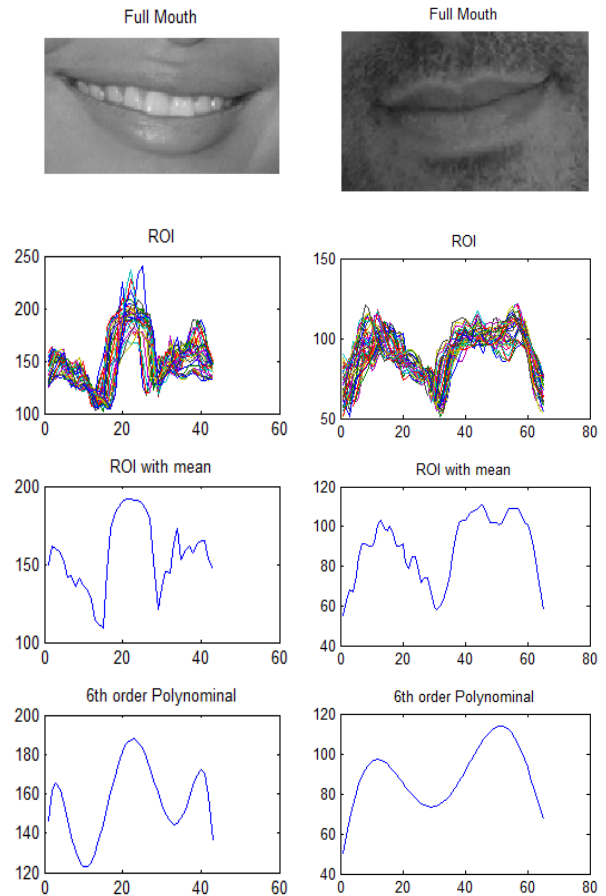


Figure 6: correct detection of a smiling mouth (left) and correct detection of a non-smiling mouth (right).