Experiments on Eigenfaces Robustness

Alexandre Lemieux and *Marc Parizeau*

Laboratoire de vision et systèmes numériques(LVSN),
Département de génie électrique et de génie informatique,
Université Laval, Ste-Foy (Qc), Canada, G1K 7P4.
E-mail: {alemieux, parizeau}@gel.ulaval.ca

Abstract

This paper is an experimental study on the robustness of the eigenfaces method for face recognition. To build a face recognition system, especially in an unconstrained surveillance system where a clear, direct, and normalized view of the face cannot be assumed, one needs to implement several image preprocessing steps like segmentation, deskewing, zooming, rotation, warping, etc., before processing the face image per se. Our aim is to determine how efficient these preprocessing steps must be in order to apply the eigenfaces method with success. The experiments are conducted on a subset of the AR-face color image database. Real images are used and altered synthetically to study the effects of 7 parameters that can be translated into corresponding preprocessing artifacts: horizontal and vertical translations, downsampling, zooming, rotation, morphing and lighting.

1 Introduction

Face recognition is an active field of research, ranging from security applications to human/computer interaction. Many techniques have been developed and huge progress has been made, yielding systems with good human detection, tracking and recognition abilities [4, 5, 7, 8].

In the context of visual surveillance, ideal conditions for person identification are rarely met, therefore the need for a robust system. To use face recognition, individuals must first be found in the scene, their face segmented as accurately as possible, and then warped into a view compatible with the recognition system requirements. This last step is critical as many errors can occur. In particular, problems with lighting conditions, low resolution, head tilt/rotation, and face features detection, are especially challenging.

One of the most widely used face recognition technique is the *eigenfaces* method, based on a Karhunen-Loève procedure [2, 6]. This technique has shown good results on various databases [8]. But, as shown in the past, its effectiveness rely heavily on the quality of the image segmentation and preprocessing. Most systems and algorithms require centered and normalized face images, thus assuming that these steps are flawless.

Our aim in this paper is to study the effects of the following parameters on the eigenfaces method: horizontal and vertical image translation, downsampling, scale, rotation, morphing, and lighting. These parameters can be linked to corresponding preprocessing artifacts. For instance, erroneous feature detection can lead to inadequate normalization, and thus to image translation, rotation, and scale. Also, the lack of image resolution is equivalent to image downsampling, and head tilt or head rotation produces some sort of image morphing.

The paper is organized as follows. Section 2 first describes the experimental protocol used to conduct our experiments. Experimental results and their analysis are then presented in Section 3. Section 4 concludes the paper.

2 Experimental protocol

Our experiments were conducted using a subset of the AR-face image database that contains over 3200 color images of 126 persons taken during two distinct sessions, with different face expressions and poses [3]. The original images are 768x564 pixels with 24 bits per pixel. From this database, 133 frontal views of men and women[1] were extracted and normalized prior to eigenfaces computation. Figure 1 shows the average image together with the first 4 eigenfaces (from left to right). Image normalization was conducted using the following procedure:

1. Eyes centers and nose tips were manually added to the database;
2. Images were rotated so that both eyes are perfectly aligned horizontally;

[1] Although the database documentation mentions only 126 different persons, we actually found 133 that belongs to the same session.

Figure 1. Average image + first 4 eigenfaces.

Figure 2. Examples of parameter effects. From left to right: original image, after downsampling (90%), after scale (-17.5%), after rotation ($20°$), after morphing (30%), after luminance changes ($\times 1.4$).

3. Images were scaled to set the distance between the eyes at 113 pixels exactly[2];
4. Images were translated so that all eyes are perfectly aligned;
5. Images were cropped to size 200x220 around a point that has the horizontal coordinate of the mid-point between the eyes, and the vertical coordinate of the average nose tip.

It is important to note that the training data set consists of only one image per individual. The reason is that we want to test the eigenfaces method using images that differ from the training images only by a controlled parameter, so that the sensitivity of the approach relative to each parameter can be quantified more precisely. In our experiments, the overall recognition is not really relevant, only the relative performances are. Moreover, in eigenfaces space, distances between images and classes will be compared with inter-class distances.

The following paragraphs present the 7 studied parameters in more details. Figure 2 illustrates the effects of 5 of these parameters (translations are not illustrated), in relation with the original image (leftmost). The figure gives images for three different persons including one with eye glasses.

Translation Up, down, left and right translations are tested over the range up to 40% of the total width or height of the test image. This represents a maximum of 80 pixels for horizontal translations, and 88 pixels for vertical translations.

Downsampling In a visual surveillance scene, depending on the distance of the subject, the resolution of the face to recognize can be much lower than the one that was used for training. Thus, an unknown face image must typically be enlarged by digital zooming. Accordingly, the downsampling parameter defines the loss of resolution in percentage of image size. For instance, downsampling by 80% means that the resolution of a 200×220 image will decrease to 40×44. This downsampling is achieved by a simple nearest neighbor algorithm in the range from 50% to 99%. Afterwards, the image is scaled back to its original size using a cubic spline interpolation algorithm. The effect of downsampling is similar to blurring.

Scale The scale parameter modifies the image size to simulate an incorrect normalization. This parameter is varied from -40% to $+40\%$ (a negative scale shrinks the image, while a positive scale corresponds to a zoom effect).

Rotation The rotation parameter rotates the image around its center counterclockwise. It varies between $-40°$ and $+40°$ (0 degree points north). It is well known that eigenfaces are not invariant to rotation [6], and that it should be corrected by the normalization transform to achieve good recognition result.

Morphing – Pseudo-rotation The morphing parameter seeks to simulate a rotation of the head around its axis. It defines an image transform that expands the left half of the face, and contracts the right half (see Figure 2, 5th column). Obviously, this parameter does not pretend to be a realistic rotation of the head (in-depth), but nevertheless it does simulate a non-linear deformation of the face that should stress the eigenfaces method in a similar fashion. A morphing of 30% expand the left half of the face by 30% and compresses the right half also by 30%. The parameter effect is studied in the range from 0 to 75%.

Luminance Many articles have investigated the effect of different lighting conditions on face images [1]. When a face is segmented in a scene, its luminance histogram can be adjusted to fit any one of several models. The luminance parameter in this experiment applies a scaling factor to the luminance channel of the image (in HLS space) which affects both the average of its histogram and its variance. The parameter varies from 0 to 1.5.

3 Experimental results

To evaluate the effect of each parameter, two different measures are reported: 1) recognition rate via a nearest-neighbor decision rule, and 2) normalized root mean square (RMS) error in eigenfaces space. In the first case, the unknown face is simply projected onto the eigenfaces space to

[2] This value corresponds to the average distance before normalization.

Table 1. Independent parameter limits.

Parameter	Limits
Morphing	$< 20\%$
Rotation	± 5 degrees
Scale	$\pm 5\%$
Translation Up-Down	$\pm 3\%$ of height
Translation Left-Right	$\pm 5\%$ of width
Downsampling	$< 90\%$

become a point. This point is then compared with the class prototypes associated with the training faces, and the unknown image is classified into the nearest-neighbor class. In the second case, a distance metric is computed from the RMS error between the projected images in eigenfaces space and their associated class prototypes. This value is finally normalized by the average minimum inter-class distance. For the original images, these measurements give respectively 100% recognition and 0 error (all eigenfaces are used).

Results are given in Figure 3 for the 7 studied parameters. Two graphs are given for each parameter: 1) the recognition rate (left scale), and 2) the normalized RMS error (right scale). Figures 3a and 3b show clearly that eigenfaces are very sensitive to translation error, especially vertically. A vertical misalignment of only 11 pixels (5%) breaks the algorithm by more than 40%. This implies that the detection of face features for normalizing the images is a very crucial step. Figure 3c shows that, on the contrary, eigenfaces are very robust to low resolution images. We were able to downsample all the way to 94% (from 200×220 to 12×13) before any loss of recognition rate, and a relatively low RMS error increase. This seems to imply that eigenfaces use mostly low frequency information for recognition. It also implies that face recognition from far away is possible, at least if the face is looking at the camera!

Scale results are shown in Figure 3d. They suggest that eigenfaces are more tolerant on slightly wrong scales than on misalignments. Moreover, the asymmetry in the curves indicate that it is preferable to oversize the faces rather than to undersize them. A possible explanation is that undersizing the faces introduces more hair or background pixels into the process (in the case of the AR-face database, this background is white however). Rotation results are shown in Figure 3e. Like for scales, eigenfaces seem to be robust enough to tolerate a reasonable amount of rotation error. Also some non negligible amount of morphing is tolerated by the algorithm, as shown in Figure 3f. This suggests, that it will be possible to recognize faces from far away, even if a direct frontal view is not possible. Figure 3g shows the effect of changes in luminance. Contrary to other parameters, it is interesting to note that the error varies linearly with a change in luminance. Table 1 summarizes the range of independent parameter values for which the eigenfaces method seems to be robust. Finally, the combined effects of all parameters are illustrated in Figure 3h, where abscissa 0 is for the original images, and interval $[-1, +1]$ represents the combined lower and upper limits of Table 1 (applied in that order). This last graph shows that the eigenfaces method is robust, even in the lack of ideal conditions.

4 Conclusion

Results have shown that, overall and up to a certain point, eigenfaces are robust over a wide range of parameters. However, they also show that, passed this point, the algorithm can breakdown sharply. This is true especially for horizontal and vertical misalignments, luminance changes, and low resolution faces (although this last parameter is robust over a very large interval). Results indicate that eigenfaces are quite robust to both low resolution images and warped images, which suggest that it can be used for face recognition even when the faces are far away from the camera, assuming that preprocessing steps can segment the faces and extract sufficient features for adequate normalization. Our interest in eigenfaces lie in an application of human video surveillance where the system not only uses face recognition to identify passerby, but also other morphological body features.

Acknowledgments This research was supported by an NSERC–Canada grant to M. Parizeau.

References

[1] P. N. Belhumeur, J. Hespanha, and D. J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE PAMI*, 19(7):711–720, July 1997.

[2] M. Kirby and L. Sirovich. Application of the karhunen-loeve procedure for the characterization of human faces. *IEEE PAMI*, 12(1):103–108, 1990.

[3] A. M. Martinez and R. Benavente. The AR-face database. Technical report, CVC Technical Report #24, June, 1998.

[4] A. M. Martinez and A. C. Kak. PCA versus LDA. *IEEE PAMI*, 23(2):228–233, 2001.

[5] B. Moghaddam, W. Wahid, and A. Pentland. Beyond eigenfaces: Probabilistic matching for face recognition. In *Proc. of Int'l Conf. on Automatic Face and Gesture Recognition (FG'98)*, pages 30–35, Nara, Japan, April 1998.

[6] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of cognitive neuroscience*, 3(1):71–86, 1991.

[7] C. R. Wren, A. Azarbayejani, T. Darrell, and A. Pentland. Pfinder: Real-time tracking of the human body. *IEEE PAMI*, 19(7):780–785, 1997.

[8] W. Zhao, R. Chellappa, A. Rosenfeld, and P. Phillips. *Face Recognition: A Literature Survey*. http://citeseer.nj.nec.com/374297.html, 2000. 67 pages.

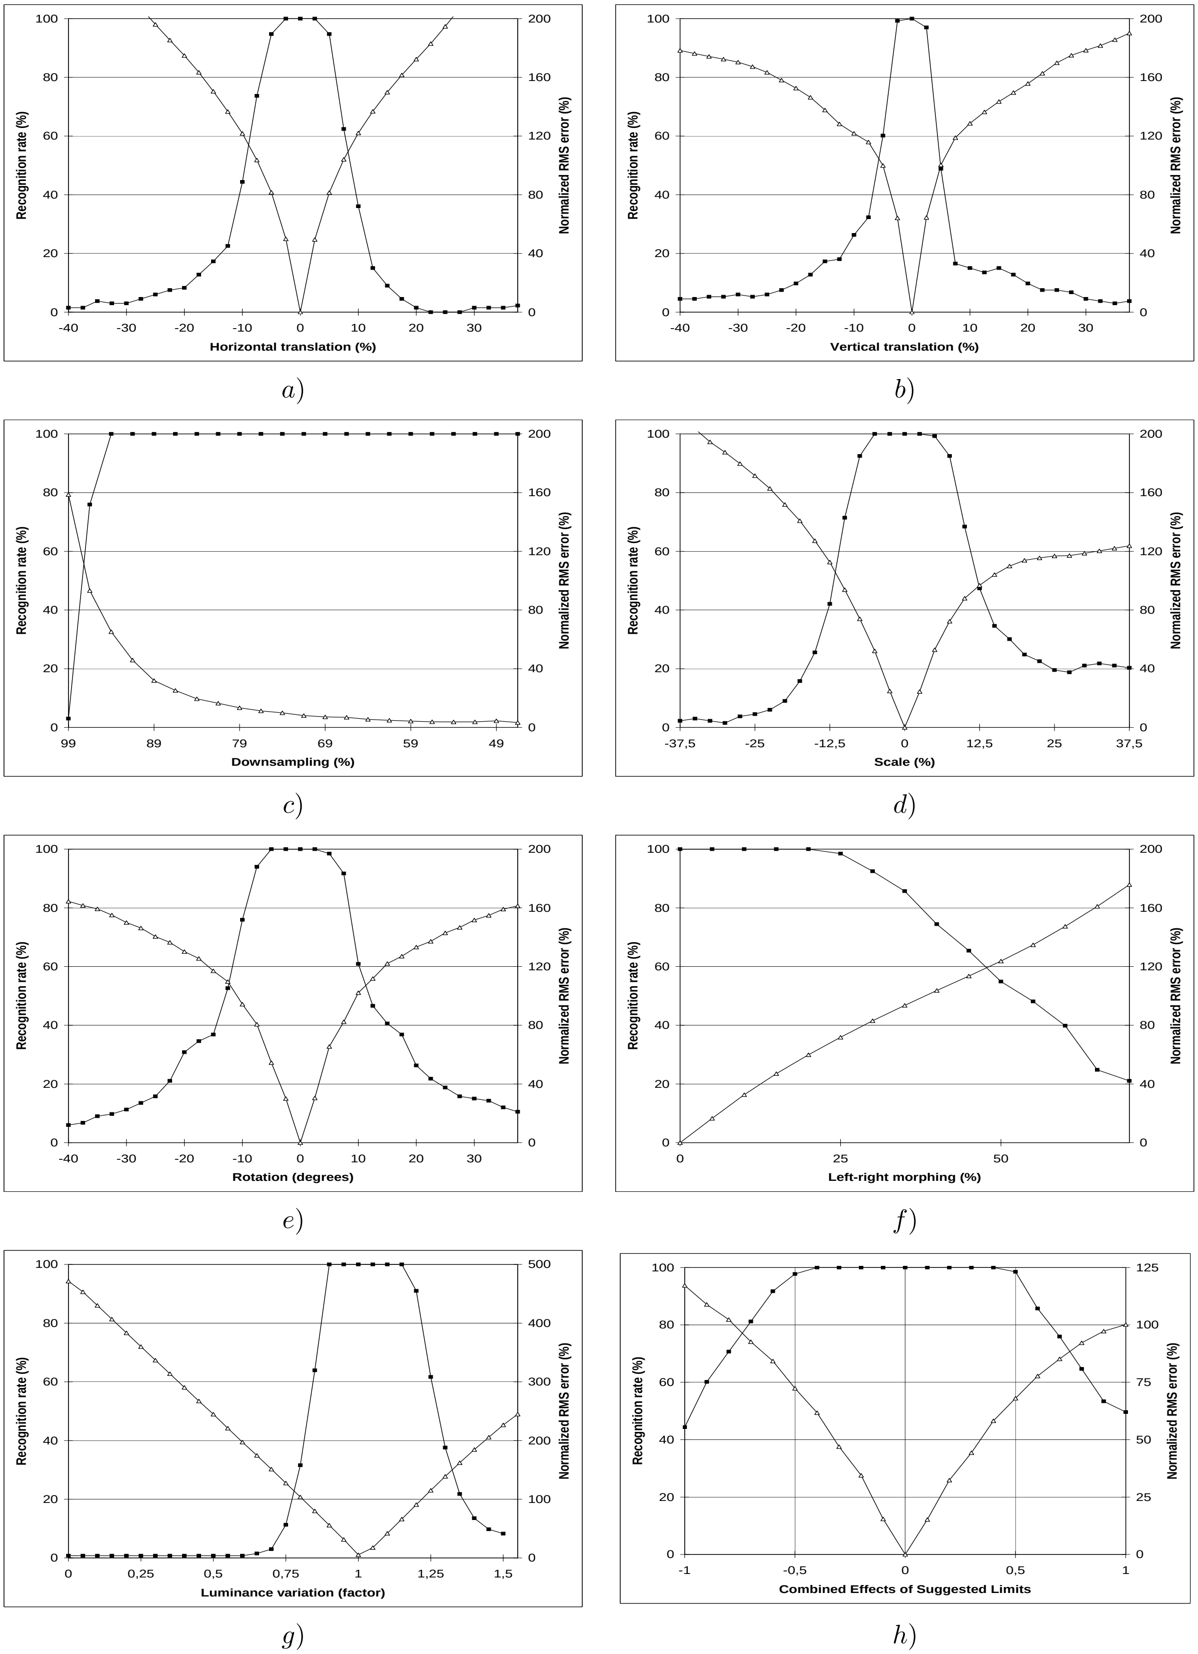

Figure 3. Effects of parameters: ■ — Recognition rates (left scale); △ — RMS error (right scale).