How Much Stress are Presidents Under?

A Study of Age Progression in Automated Face Recognition

Rachel Cindric
MERIT Program
University of Maryland
College Park

Narayanan Ramanathan
ECE Department
University of Maryland
College Park

Rama Chellappa
ECE Department
University of Maryland
College Park

Abstract

One of the major challenges faced by automatic face recognition is caused by aging. An algorithm has recently been developed that can not only recognize a face that has aged, but can also tell how much the face has aged. We collected several photographs of U.S. presidents before and after they were in office and sent them through the algorithm. The results will show how much each president appeared to age while he was in office, whether it was more than he should have or less.

1. Introduction

Automated recognition of human faces has received increased attention in recent years. Much research has been devoted to solving the problems of occlusion, expression, illumination, and pose. Zhao et al. [5] gives a detailed review of the recent progress in automated face recognition.

Age progression has not been researched as much as the other problems, but it is just as important for face recognition. A significant age gap can cause degradation in automated face recognition, especially if an image of a younger person is being used. Up until about the age of 20, aging is seen mostly in the changing of the shape of the skull. After that, aging can be seen primarily in skin changes, most notably through wrinkling.

.

1.1 Background of Aging Research

Kwon and Lobo [1] used ratios between the eyes, nose, mouth, chin, and top of the head to distinguish between children and adults. They also dropped snakelets onto the images in order to use wrinkles to distinguish between young adults and old adults.

Lanitis et al. [2] developed a way to model aging variations and use them to estimate aging in a person and also to simulate aging in an unseen person. They also used these methods to create an age-robust face recognition system.

1.2 Purpose Statement

An algorithm has recently been developed by Ramanathan and Chellappa [4] that can determine the age difference between two face images. This raises an interesting question: Does a face visibly age more than normal if a person is under a lot of stress? We decided to test this idea using images of U.S. presidents, who are undoubtedly under a lot more scrutiny and stress than the average person.

We collected 61 images of 11 different presidents from before and after they were in office. Each president had at least one before and after photo, with some presidents having several of each. Out of these 61 images, only the before and after pairs with the best illumination and mostly frontal pose were sent through the algorithm.

In section 2, a detailed explanation is given of the steps the images go through before they are entered in the algorithm. The algorithm itself is also explained in this section. Section 3 gives the results of the experiment, while explaining in which cases the algorithm worked and in which cases it didn't. Conclusions are made and ideas for improvement are given in section 4.

2. The Process

2.1 Face Cropping

Before the images can be entered into the age-difference algorithm they must all be cropped down to the same size, 207x180 pixels. Since a face is roughly symmetric, we used the location of the eyes to crop all the faces, so that the eyes of every face were lined up. This was done using a MATLAB program where the user clicks on the eyes of the face and the image is cropped accordingly. The faces were also placed in ellipse masks to eliminate the different backgrounds, as shown in figure 1. Any images that were originally in color were also converted to gray scale when they were cropped.

It is very important that each pair of before and after faces is the same size. If a before face is cropped differently than its corresponding after face, then the age-difference algorithm will possibly view that difference as an effect of aging. Because of this, some of the faces had to be cropped numerous times, until they were close enough in size to the other face in the pair.

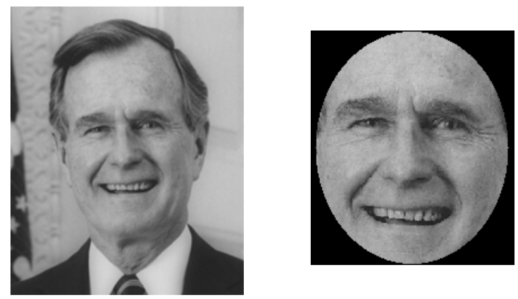

Figure 1: An original image of President Bush and its corresponding cropped and masked face

2.2 Half-faces

It has been shown that using half of a face instead of the whole face can eliminate some pose and illumination problems and therefore improve results [4].

We took each face image and decided which half had better pose and illumination. Each full face is made up of 207x180 pixels with each pixel occupying a place in a matrix. If the left half of the face is chosen, the first 90 rows of the matrix are kept the same while the last 90 are eliminated. If the right half is chosen, the last 90 rows of the matrix are flipped to make them the first 90 rows. This way all of the half faces are facing the same way, no matter which side was chosen as shown in figure 2.

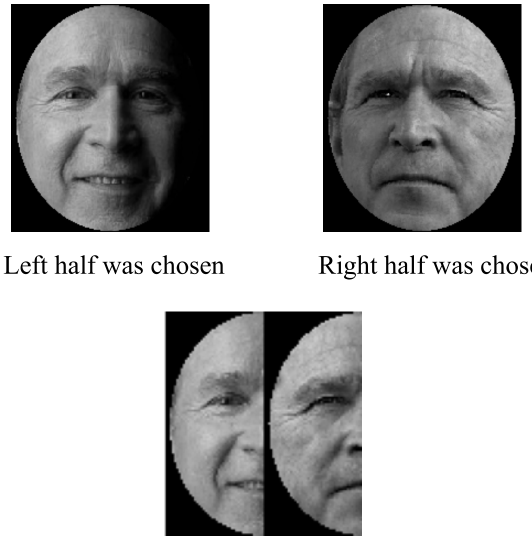

Figure 2: Two face images and the half -faces that were chosen from each

2.3 Age Categorization

The age-difference algorithm [4] takes two face images and returns one of 4 age categories: 1-2 yrs, 3-4 yrs, 5-7 yrs, or 8-9 yrs. It does this in two steps. First it uses intra-personal and extra-personal differences to determine whether the two face images are the same person. Then, if they are the same person, intra-age differences are used to determine how many years apart the images are.

The database for the algorithm consists of 465 different pairs of faces as shown in table 1. To create the intra-personal space, Ω_I, 200 of these pairs were selected and their intra-personal differences, **x**, were found using

$$\mathbf{x}_i = I_{i1} - I_{i2} \tag{1}$$

where I_{i1} is the first face in the ith pair and I_{i2} is the second face in the ith pair.

Table 1: Database of images for age difference algorithm

Age Difference	1-2 yrs	3-4 yrs	5-7 yrs	8-9 yrs
# of pairs	165	104	81	115

To create the extra-personal space, Ω_E, two faces at a time from different pairs were chosen out of 200 pairs of images. The extra-personal differences, **z**, were found using

$$\mathbf{z}_i = I_{i1} - I_{j2},\ j \neq i \tag{2}$$

where I_{j2} is the second face in the jth pair.

Both the intra-personal and extra-personal spaces can be broken into two complementary subspaces. These subspaces are the feature space F and the orthogonal complement space $\bar{F}$.

The likelihood function for the intra-personal data is:

$$P(\mathbf{x}|\Omega_I) = \frac{\exp(-\frac{1}{2}(\mathbf{x}-\bar{\mathbf{x}})^T\Sigma^{-1}(\mathbf{x}-\bar{\mathbf{x}}))}{(2\pi)^{N/2}|\Sigma|^{1/2}} \tag{3}$$

Moghaddam and Pentland [3] suggest that instead of evaluating this problem explicitly it is easier to estimate the likelihood function as:

$$\hat{P}(\mathrm{x}|\Omega_I) = \frac{\exp(-\frac{1}{2}\sum_{i=1}^{k}\frac{y_i^2}{\lambda_i})}{(2\pi)^{k/2}\prod_{i=1}^{k}\lambda_i^{1/2}} \cdot \frac{\exp(-\frac{\varepsilon^2(x)}{2\rho})}{(2\pi\rho)^{(N-M)/2}} \tag{4}$$

where y_i are the principal components, λ_i are the eigenvalues, ε2(**x**) is the PCA reconstruction error, and ρ is estimated by extrapolation of the cubic spline fit on the eigenvalues.

The likelihood function for the extra-personal space can be estimated as:

$$\hat{P}(\mathbf{z}|\Omega_E) = P(\mathbf{y}|\Theta^*) \cdot \hat{P}_{\bar{F}}(\mathbf{z}|\Omega_E) \tag{5}$$

where

$$P(\mathbf{y}|\Theta) = \sum_{i=1}^{N_c} w_i N(\mathbf{y}; \mu_i, \Sigma_i)$$

$$\Theta^* = argmax \left[\prod_{i=1}^{M} P(\mathbf{y}_i|\Theta) \right] \tag{6}$$

$N(y; \mu_i, \Sigma_i)$ is Gaussian and w_i are the mixing parameters such that $\sum_{i=1}^{N_c} w_i = 1$. This problem can be solved using the Expectation-Maximization algorithm.

Using the Bayes rule, the *a posteriori* probability can be found:

$$P(\Omega_I|\mathbf{x}) = \frac{P(\mathbf{x}|\Omega_I)P(\Omega_I)}{P(\mathbf{x}|\Omega_I)P(\Omega_I) + P(\mathbf{x}|\Omega_E)P(\Omega_E)} \tag{7}$$

Using this probability, a new pair of face images can be classified as intra-personal or extra-personal. $P(\Omega_I)$ and $P(\Omega_E)$ are set equal, and if $P(\Omega_I|\mathbf{x}) > ½$, both images in the pair are of the same person.

In a similar fashion the age category for the new pair of faces can be found. One intra-personal space was built for each of the four age categories using 50 images from each category. These spaces are represented by Ω_1, Ω_2, Ω_3, and Ω_4. The a posteriori probability for each space is found using:

$$P(\Omega_i|\mathbf{x}) = \frac{P(\mathbf{x}|\Omega_i)P(\Omega_i)}{\sum_{j=1}^{4} P(\mathbf{x}|\Omega_j)P(\Omega_j)} \tag{8}$$

where **x** is the difference from the new pair of faces that was classified as intra-personal. If $P(\Omega_i|\mathbf{x}) > P(\Omega_j|\mathbf{x})$ for all $i \neq j$, $i,j = 1,2,3,4$, then the new pair of faces belongs to the age category i.

3. Experimental Results

Using half-faces instead of the full face can help eliminate problems with pose and illumination. Figure 3 shows a full face and two pairs of half faces for President Nixon. Both pairs are from the same before and after photos, with the first pair using the right side of the before-face and the second pair using the left side of the before-face. A third pair of half-faces was also sent through the algorithm and returned 5-7 yrs for the age category, which matched the category returned by the second pair. The right side of the before-face had a different pose and did not match up well with the second face. So using the left side of the face improved the ability of the algorithm to correctly identify the aging.

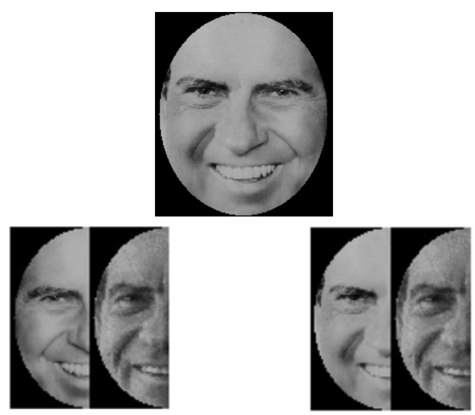

8-9 yrs   5-7 yrs

Figure 3: Full and half-faces of President Nixon. The first pair includes the right half of the whole face and the second pair includes the left half

Figure 4: Three pairs of half-faces for President Reagan that all returned the same age category of 3-4 yrs

For a few of the presidents, more than one pair of images was used. In the case of President Reagan, three different image pairs from the same years were sent through the algorithm, as seen in figure 4. The after-face was the same for each pair. Each of these pairs returned the same age category, even though the illumination in the before-faces was fairly different. In this case the algorithm was able to handle the illumination differences quite well.

Five pairs of images from the same years were sent through for President Kennedy as shown in figure 5. These pairs did not all return the same age category. The illumination and pose were very different for the images in some of the pairs. Also, some of the faces had different intensities and quality. There were too many differences in these cases for the algorithm to work accurately. The pairs of images that returned 5-7 yrs seem to be matched much better than the ones that returned 3-4 yrs, so we assume that the correct age category is 5-7 yrs.

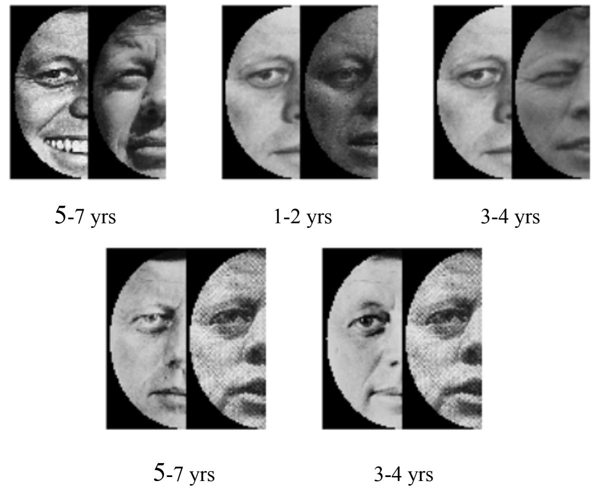

Figure 5: Five pairs of half-faces for President Kennedy that returned several different age categories

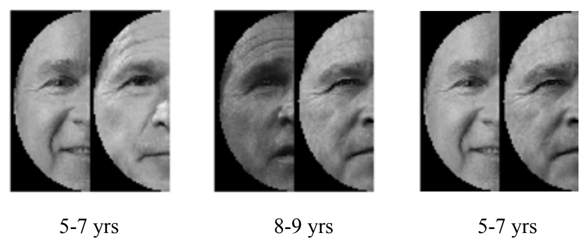

Figure 6: Three pairs of half-faces for President W. Bush. The center pair returned erroneous results

For President G. W. Bush, three different pairs of images were used. Two of the pairs were very similar in illumination, while the other one was significantly different as seen in figure 6. An age category of 8-9 years is returned because the algorithm mistakes the difference in illumination as a sign of aging. The results of the poorly illuminated pair can be tossed out and it can be assumed that the correct age category is 5-7 yrs.

The age categories for all of the presidents as well as the actual age difference between the faces are shown in table 2.

	Age Difference	Age Category Returned
W. Bush	4 yrs	5-7
Clinton	4 yrs	5-7
Bush	4 yrs	3-4
Reagan 2nd	6 yrs (‘87-’93)	5-7
Reagan 1st	6 yrs (’81-’87)	3-4
Carter	4 yrs	3-4
Ford	3 yrs	1-2
Nixon	6 yrs	8-9
Johnson	5 yrs	5-7
Kennedy	3 yrs	5-7
Eisenhower	4 yrs	3-4
Roosevelt 2nd	6 yrs (’38-’44)	5-7
Roosevelt 1st	5 yrs (’33-’38)	3-4

Table 2: Actual age difference and age category returned by the algorithm for each president

4. Conclusions

From looking at table 2, it can be seen that about half of the presidents appeared to have aged about the number of years that they should have. Of the remaining presidents, about half aged more than normal and half aged less then normal.

An interesting extension of this research would be to look at the presidents who appeared to age more or less than they should have and see the different issues that they had to deal with while they were in office. Is it possible that the ones who appeared to age more had a much harder presidency, while the ones who aged less had it fairly easy?

Improvements could be made in the results by using photos that have better pose and illumination. Although it is quite hard to find face images that have similar pose and illumination, it is possible to do so with the right resources.

References

[1] Y. Kwon and N. Lobo, “Age Classification from Facial Images,” *Computer Vision and Image Understanding*, 74(1):1-21, April 1999.

[2] A. Lanitis, C. Taylor, and T. Cootes, “Toward Automatic Simulation of Aging Effects on Face Images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):442-455, April 2002.

[3] B. Moghaddam and A. Pentland, “Probabilistic Visual Learning for Object Representation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):696-710, July 1997.

[4] N. Ramanathan and R. Chellappa, “Face Verification across Age Progression,” *IEEE Computer Vision and Pattern Recognition*, 2:462-469 June 2005.

[5] W. Zhao, R. Chellappa, P. Phillips, and A. Rosenfeld, “Face Recognition: A Literature Survey,” *ACM Computing Surveys*, 35(4):399-458, December 2003.