Facial feature extraction and pose determination

Athanasios Nikolaidis Ioannis Pitas *

Department of Informatics

Aristotle University of Thessaloniki

Box 451, Thessaloniki 540 06, GREECE

E-mail: {nikola, pitas}@zeus.csd.auth.gr

*Corresponding Author

Summary

The present paper describes a set of methods for the extraction of facial features as well as for the determination of the gaze direction. The ultimate goal of the proposed approach is to identify a sufficient set of features and distances on them aiming at a unique description of the face structure. The extracted feature set and feature distances can provide a robust representation of a facial image. This representation is suitable for identifying a person within a large image database using a face recognition system. Another suitable application would be to combine the extracted feature set and the gaze direction information in order to track the movement of a face in a sequence of facial images. Eyebrows, eyes, nostrils, mouth, cheeks and chin are considered as useful features. The candidates for eyes, nostrils and mouth are determined by searching for minima and maxima in the $x-$ and $y-$projections of the greylevel relief. An unsupervised clustering method is used in order to select the dominant candidate. The candidates for eyebrows are determined by adapting a proper binary template to an area restricted by the position of the eyes. The candidates for cheeks and chin are determined by performing an adaptive Hough transform on a relevant subimage defined according to the size and position of the ellipse that describes the largest connected component of the image. In the absense of marginal features of the face, such as the cheeks and the chin, a dynamic programming technique based on estimating this ellipse that represents the main face region is applied in order to acquire an estimate of the inner face contour. Finally, the direction of gaze is determined by using the symmetric properties of certain facial features and by applying rules of projective geometry. The algorithms presented were tested on the M2VTS multimodal face database.

Abstract

A combined approach for facial feature extraction and determination of gaze direction is proposed that employs some improved variations of the adaptive Hough transform for curve detection, minima analysis of feature candidates, template matching for inner facial feature localization, active contour models for inner face contour detection and projective geometry properties for accurate pose determination. The aim is to provide a sufficient set of features for further use in a face recognition or face tracking system.

Keywords

Minima analysis, Adaptive Hough transform, Template matching, Active contour models, Projective geometry, Biometric properties.

1 Introduction

One of the key problems in building automated systems that perform face recognition tasks is face detection and facial feature extraction. Many algorithms have been proposed for face detection in still images that are based on texture, depth, shape and colour information. For example, face localization can be based on the observation that human faces are characterized by their oval shape and skin-colour[(1)]. Another very attractive approach for face detection relies on multiresolution images (also known as mosaic images) by attempting to detect a facial region at a coarse resolution and, subsequently, to validate the outcome by detecting facial features at the next resolution [(2,3)]. A critical survey on face recognition can be found in (4).

The major difficulties encountered in face recognition are due to variations in luminance, facial expressions, visual angles and other potential features such as glasses, beard, etc. This leads to a need for employing several rules in the algorithms that are used to tackle these problems.

Several techniques for face recognition have been proposed during the last years. They can be divided into two broad categories: techniques based on grey level template matching and techniques based on the computation of certain geometrical features. An interesting comparison between the two categories can be found in (5). Experimental research through years proved that each category covers techniques that perform better on the localization of certain facial features only. This implies that for every selected feature a fusion of methods of both categories should provide more stable results than the corresponding method alone.

Following similar reasoning, in this paper, we propose the use of complementary techniques that are based on geometrical shape parameterization, template matching and dynamic deformation of active contours, aiming at providing a sufficient set of features that may be later used for the recognition of a particular face. The success of a specific technique depends on the nature of the extracted feature. In the present paper, the features under consideration are the eyebrows, the eyes, the nostrils, the mouth, the cheeks and the chin. The following methods are employed throughout the paper: i) the extraction of eyes, nostrils and mouth based on minima analysis of the x- and y-greylevel reliefs (section 2.1), ii) the extraction of cheeks and chin by performing an Adaptive Hough Transform (AHT)[(6)] on the relevant subimage defined according to the ellipse containing the main connected component of the image (section 2.2), iii) the extraction of upper eyebrow edges using a template-based technique that adapts a proper binary mask to an area restricted by the position of the eyes (section 2.3), iv) the accurate extraction of the face contour by means of a dynamic deformation of active contours (e.g. snakes), compensating for the lack of parametric information needed in the AHT (section 3), and v) the determination of the gaze direction by exploiting the face symmetry properties (section 4). Certain improvements in the proposed implementations of both the AHT and the snakes techniques, compared to some previous works in the same field, are discussed throughout the corresponding sections. Some experimental results that demonstrate the success of the proposed approach are presented in section 5. Finally, certain conclusions are drawn in section 6.

2 Facial feature extraction

As described in (1), skin-like regions in an image can be discriminated by representing the image in the HSV (Hue-Saturation-Value) colour space and by choosing proper thresholds for the values of the hue and saturation parameters. Other similar colour spaces that reside to the way the human vision interprets colour information, like HSI (Hue-Saturation-Intensity) or HLS (Hue-Luminance-Saturation), could also be used. Considering the oval-like shape of a face in general, it is convenient to search for the connected components of the skin-like regions using a region growing algorithm and to fit an ellipse to every connected component of nearly elliptical shape. This is a basic preprocessing step, because it restricts the search area for the facial feature detection.

2.1 Eyes, nostrils and mouth localization

The extraction of features like eyes, nostrils and mouth relies on the fact that they correspond to low intensity regions of the face. In the case of the eyes, this is due to the colour of the pupils and the eye-sockets. As proposed in (1), first morphological operations are applied on the image in order to enhance the dark regions. After a normalization of the orientation of the face candidate, the x- and y-projections of the greylevel relief are computed. This helps us to search for facial features only along the horizontal direction.

The y-projection of the topographic greylevel relief is first computed by averaging the pixel intensities of every row. Significant minima are determined based on the gradient of each minimum to its neighbour maxima. This provides us with the possible vertical positions of the features. Afterwards the x-reliefs of the rows corresponding to the detected minima are computed, and minima analysis follows once again in order to find the potential horizontal positions of the features.

After constructing the lists of significant minima and maxima, candidate features are determined based on several metrics, e.g. the relative position inside the head, the significance of maximum between minima, the distance between minima, the similarity of greylevel values, and the ratio of distance between minima to head width. Finally, the best complete set of features is determined based on the vertical symmetry of the set, the distances between facial features and the assessment of each facial feature candidate, based on fuzzy set theory.

2.2 Parameteric feature extraction using AHT

In order to select an efficient set of distances, a parameterized representation of some facial features is a good approach towards the solution of the problem under study. Such an assumption holds for certain features that can be described in a realistic way by means of a geometric curve. The cheeks can be approximated reasonably well by nearly vertical lines, and the chin by an upward parabola. A straightforward method to detect geometric curves in an image is the Hough transform or one of its variations [(7)]. The use of an adaptive Hough Transform is proposed in (8). It is a variant of the Hough transform [(6)], which utilizes a small fixed size accumulator and follows a "coarse to fine" accumulation until a desired parameter precision is reached. This variation requires a reduced storage space compared to the standard transform due to the small accumulator size. Accordingly, it provides a better performance than the standard transform due to the reduced number of computations. Cheek border detection is based on the fact that a straight line is represented by a linear equation of the form:

$$x = my + c \tag{1}$$

where m is the slope and c is the intercept of the line. Thus the accumulator array is two-dimensional, with each cell corresponding to a certain pair (m, c) of discrete values of the slope and of the intercept of the line. The inherent problem in equation (1) that it cannot describe horizontal lines, is not a concern for cheek line detection. Nearly vertical lines can easily be represented for $m \simeq 0$. The cheek border corresponds to the dominant vertical line in the relevant subimage.

Accordingly, the algorithm should search only for one line in the subimage. The main idea of the method is to focus on the parameter space on areas of large accumulation by performing iterations of the basic Hough algorithm and redefining the limits of the parameter range accordingly. If we let $a_i, 1 \leq i \leq n$ be the parameters representing the curve under consideration (in this case $n = 2, a_1 = m, a_2 = c$), then the n-dimensional search space in iteration i is

$$S^{(i)} = L_1^{(i)} \times L_2^{(i)} \times \ldots \times L_n^{(i)} \tag{2}$$

where $L_j^{(i)}$ is an interval (U_j, V_j) of possible values of parameter a_j.

The relevant subimage on which to apply the algorithm should be adapted to our needs. In (8), the relevant subimage is defined according to the distance between the eyes and the width of the mouth. In our case, the relevant subimage is defined as a region of the largest best-fit ellipse found in the image, which acts as a good initial description of the face prior to feature detection and provides more stable results because it considers the symmetric properties of the specific face in more detail. The best-fit ellipse is assumed to contain the main facial region. The success of a search inside the ellipse is ensured by the fact that the ellipse possesses the same vertical symmetry with its enclosed face region. After experimentation and by taking into account some biometric analogies of the human face, it has been found that AHT should perform on a well chosen region of interest for best performance.

Let a and b be the magnitudes of the major and minor semiaxes and (x_0, y_0) be the center of the best-fit ellipse. The boundaries of a suitable subimage for our purposes, should be of the form:

$$x_{min} = x_0 + p_{high} \cdot a, \tag{3}$$

$$x_{max} = x_0 + p_{low} \cdot a. \tag{4}$$

$$y_{min} = y_0 + p_{left} \cdot b, \tag{5}$$

$$y_{max} = y_0 + p_{right} \cdot b. \tag{6}$$

where $p_{high}, p_{low}, p_{left}, p_{right}$ denote appropriate percentages of the semiaxes magnitudes.

In the case of not perfectly frontal face images, the best-fit ellipse is rotated before the search procedure so that the main head axis is parallel to the x-axis[(9)].

The chin can be described by the equation of an upward parabola:

$$x - x_v = g(y - y_v)^2 \tag{7}$$

where (x_v, y_v) is the vertex of the parabola and g is the parameter which controls how fast the parabola opens outwards. In this case, the accumulator array is three-dimensional ($n = 3$) having parameters $a_1 = g, a_2 = x_v, a_3 = y_v$. The relevant subimage used for searching is defined as a certain region of the best-fit ellipse as well. By exploiting the inherent symmetry, we prevent Hough transform from extracting some other features resembling to a parabola, like the lower edge of the lips.

A compensation for rotations along the z-axis is again performed prior to the search procedure.

2.3 Eyebrows extraction using template matching

Unlike the cheeks and the chin, the eyebrows cannot easily be described by means of a geometric curve, because of their variations among various people. One way to cope with this problem is to consider only the upper edge of the arch of the eyebrow for detection. This edge constitutes a transition in the vertical direction, from a completely smooth area (the forehead) to an elongated approximately horizontal dark region. This is illustrated in Figure 1. A template matching technique is more appropriate in this case. A prototype block has to be defined explicitly, simulating the form of the upper edge of an eyebrow, because no *a priori* information exists concerning the approximate position of the correct block (e.g. from a previous frame). Let P be the predefined block (acting as a mask) and B_i be the real block of the thresholded gradient of the image under consideration, where $B_i \in S$ with S being the set of all blocks centered on the search area. The optimal block B_{opt} has maximal correlation to the prototype block. This is expressed by:

$$B_{opt} = \{B_j \in S \mid \sum P \cdot B_j = \max_{B_i \in S} \sum P \cdot B_i\} \tag{8}$$

where the sums are considered over the binary values of the corresponding points of the prototype block and the real block. The size of the search block should be normalized according to the scaling factor of the face in the image. A straightforward way to normalize the search process is to define the dimensions of the block based on the distance between the eyes. The optimal block is expected to contain the upper edge of the eyebrow and can be used to provide a potential facial feature.

3 Snake deformation for face contour definition

In many cases, problems arise when there is no sufficient edge information for acceptable performance of the AHT algorithm for finding the cheeks or the chin. This is usually due to the poor lighting conditions. The active contour models [10] were introduced in the field of object tracking and proves to be a helpful technique to overcome this problem.

In some previous work [11], an open convex hull that is derived from the gradient image is used as an initialisation of the active contour (snake) that describes the face contour. In our case, the best-fit ellipse is a much more accurate initial estimate for the inner contour of the face, because it is produced based on the skin-colour segmentation. The objective of the method is to consider the best-fit ellipse surrounding the face as an active contour (snake) which can be deformed so that its energy is minimized or, equivalently, its optimal position is determined.

During our experiments it was observed that the best-fit ellipse usually covers an area which is wider than the main oval region of the face because the hair was also included in the main connected component. This implies that it is difficult to reach an acceptable approximation of the face contour when applying a local evolving deformation approach on the initial ellipse. Instead of using the best-fit ellipse as the initial snake, we additionally define another ellipse of half-size axes from the original. The snake elements are now constrained to lie on certain lines joining points between the two ellipses. The points (snaxels) were chosen such that successive snaxels have approximately equal Euclidean distances.

Several approaches to active contour modeling have been proposed [12]–[14] that either employ a greedy algorithm or a dynamic programming technique. The greedy algorithm tends to perform faster but has the drawback of converging to solutions that may correspond to local minima. In our work we chose to implement a dynamic programming algorithm that provides us with the required precision for the definition of the face contour, at the cost of a somewhat slower performance. The energy of each snake point v_i is given by:

$$E_{i-1}(v_{i-1}, v_i, v_{i+1}) = \lambda_i E_{int}(v_{i-1}, v_i, v_{i+1}) + (1 - \lambda_i) E_{ext}(v_i) \tag{9}$$

where $\lambda_i \in [0,1]$ is a regularisation parameter, E_{int} denotes the internal energy of the snake, which affects the continuity and curvature of the snake, and E_{ext} denotes the external energy, which refers to the forces that pull the snake towards significant image features. In order to ensure the smoothness of the resulting contour, the internal energy term is considered to involve relations between three adjacent pixels, and thus can be written as:

$$E_{int}(v_i) = \alpha \cdot \left(|v_i - v_{i-1}|^2 + |v_i - v_{i+1}|^2\right) + \beta \cdot |v_{i-1} - 2v_i + v_{i+1}|^2 \tag{10}$$

where α and β are normalization parameters that depend on the dimensions of the specific image.

The external energy comprises of two terms. The first one corresponds to the magnitude of the output of a gradient operator (e.g. Sobel operator), and helps attract the snake to edges on the image. This is the main external energy term over regions where a transition from skin colour to background colour is dominant. The other external energy term refers to the degree of change of the Euclidean distance between adjacent pairs of candidate snaxels. This is to ensure that the smoothness of the contour is retained in regions where there is poor contrast between the face region and the background. Thus the external energy component of the snake energy can be written as:

$$E_{ext}(v_i) = E_{image} + E_{dist} \tag{11}$$

with the two terms being:

$$E_{image} = -\left|\nabla I(x,y)\right| \tag{12}$$

where $I(x,y)$ denotes the grey level function of the original image and:

$$E_{dist} = k \cdot \left|1 - \frac{d(v_{i-1}, v_i)}{d(v_i, v_{i+1})}\right| \tag{13}$$

k is a normalizing factor that depends, again, on the image dimensions, and $d(\cdot,\cdot)$ denotes the Euclidean distance between adjacent (candidate) snaxels.

The result of the above procedure is an inner contour of the face region, that excludes features attributed to hair and ears and provides a region of interest very well located for searching the inner facial features.

4 Determination of gaze direction

A facial feature extraction is related to the determination of gaze direction. Instead of using the midcenter of the eyes and the projections of the face contour on the eye horizontal as in (15), we base our calculations on the estimates of the positions of the eyes and the mouth. Once these are correctly detected, and provided that the tilt about the $z-$ axis is insignificant, a good estimate of an angle uniquely defining the gaze direction can be obtained.

Once the correct position of the eyes and mouth are determined, face symmetry properties can be exploited in order to determine the angle between the plane defined by these three facial features and the image plane. The three facial features determine an isosceles triangle. If we consider one side of the triangle to lie on the image plane, it is easy to compute a good estimate of the gaze direction angle. This is accomplished by means of trigonometric functions of angles between the lines joining the selected facial features, as shown in Figure 2. In particular, let α be the side of the triangle that lies on the image plane and connects one of the eyes and the center of the mouth. Let β be the orthogonal projection of the other side of equal size onto the image plane. This is displayed in Figure 2, where the projection of the isosceles triangle is represented by continuous lines and is denoted as ABC, whereas the isosceles triangle ABE is represented by dotted lines. This means that, if the person in the figure is staring straight towards the observer, the projected triangle should coincide with the isosceles one. The dashed lines denote a triangle CDE that lies on a plane that is perpendicular to both the image plane and the plane of the isosceles triangle. If we let θ be the angle between lines α and β, and ϕ be the angle between the image plane and the isosceles triangle plane, then we have:

$$cos\phi = \frac{|CD|}{|DE|} \tag{14}$$

from the triangle CDE, where $|\cdot|$ denotes the length of the corresponding line. From the triangle ACD we have:

$$|CD| = |AC| sin\theta \tag{15}$$

$$|AD| = |AC| cos\theta \tag{16}$$

From the triangle ADE we have:

$$|AD|^2 + |DE|^2 = |AE|^2 \tag{17}$$

and because triangle ABE is isosceles:

$$|AE|^2 = |AB|^2 \tag{18}$$

we obtain:

$$|DE| = \sqrt{|AB|^2 - |AD|^2} = \sqrt{|AB|^2 - |AC|^2 cos^2\theta} \tag{19}$$

and thus the angle ϕ can be computed using the following formula:

$$cos\phi = \frac{|AC|sin\theta}{\sqrt{|AB|^2 - |AC|^2 cos^2\theta}} \tag{20}$$

A rule-of-thumb to determine the gaze direction is to look at the isosceles triangle as it is projected onto the image plane: the biggest of the eye-to-mouth distances reveals the direction at which the person is staring. The success of the method depends, as is expected, on the correct detection of the eyes and mouth. This implies that the most accurate results are obtained when we have a precise estimation of the center of the mouth and of the eye pupils.

5 Experimental Results

The above proposed methods were tested on a set of 37 images of the M2VTS project multimodal face database. This set comprises of the best-frontal shots for each person. The choice of the specific shots among all the others was based on an algorithm that exploits the vertical symmetry axis of the face (17).

Figure 3 shows some examples of sets of correct feature results. The best-fit ellipse as well as the corresponding blocks where the cheeks and the chin were detected are indicated. The values for the parameters p_{low}, p_{high}, p_{left} and p_{right} are shown in Figure 1. The selection of these areas of the ellipse ensures that the edges found in these subimages contain part of the cheek or the chin and do not contain other significant features, like the ears. A small accumulator of size 9×9 or $6 \times 6 \times 6$ is used in the AHT algorithm for cheek or chin detection, respectively. A line segment is assumed to be interesting, if its length is at least 10 pixels in the edge image obtained by applying a Sobel edge operator (17) (of size 3×3) on the original image (of size 350×286 for the image database used in the experiments). Table 1 shows the results of extraction of cheeks, chin and eyebrows for the set of images mentioned above. A set of features was extracted in all cases. The assesment of the results was based on visual inspection and comparison to the results that one would expect based on biometrical rules.

Problems have been encountered in cheek deteection when the false symmetry of the ellipse leads to bad definition of the relevant subimage, and, thus, to an erroneous extraction of some other features considered as predominant. An example where hair is extracted instead of cheeks is shown in Figure 4(a). Similar problems cause the AHT to fail in some cases of chin detection. The inability of the edge operator to detect weak edges caused by bad luminance is more obvious here. An example of wrong chin extraction is given in Figure 4(b).

The correct extraction of the eyebrows depends on the detection of the position of the eyes, as one would expect. In our experiments the prototype was chosen to be of height equal to 0.125 of the distance between the eyes and of width equal to 0.5 of the same distance. If at least one eye is detected at a wrong position, and especially when the distance between the eyes is bigger than the actual one, the size of the block to be matched may be incorrect. This may lead to the extraction of some feature other than the desired one. Even when eyes are correctly extracted, a problem may arise if hair covers the forehead and the eyebrows, as it is shown in Figure 4(c). The percentage of correct detection of the chin in Table 1 is rather low because of the lack of sufficient edge information in that region of the face.

In the case of images that suffer from a lack of edge content, the dynamic programming algorithm succeeds in providing additional information for the definition of search space for inner facial features. This is illustrated in Figure 5. Figures 5(a), 5(b) and 5(c) show the improvement that is achieved in the definition of such features as cheeks and chin compared to Figures 4(a), 4(b) and 4(c), respectively. Both the inner (best-fit) and the outer ellipses are shown in these figures. The value of the regularisation parameter λ_i of equation (9) was fixed at 0.3 for all of the snaxels. The method fails only in extreme cases, when the main connected component of the image contains features that deform dramatically the actual face region (e.g., when no clear separating line between the hair and the forehead exists). Some other good results of the application of this method are shown in Figures 5(d), 5(e), 5(f).

Finally, the gaze direction determination technique led to correct results in all cases where eyes and mouth were extracted correctly. Figure 6(a) shows an example of a person looking to the left at an angle of 18^0 compared to a frontal picture, Figure 6(b) shows a person looking to the left at an angle of 8^0, and Figure 6(c) shows a person looking to the right at an angle of 19^0. The white triangles in these pictures correspond to the projection of the isosceles triangle on the image plane.

The average computation time for locating the full set of facial features for each person, using the described methods, was 13.6 seconds on a Silicon Graphics Indy workstation with a R4400 processor at 200 MHz, a R4000 floating point coprocessor and 64 Mbytes of main memory.

The results were good, considering the fact that the shots have not been really frontal in some cases, and that the lighting conditions did not have a uniform effect all over the image.

6 Conclusions

A combination of methods based on adaptive Hough transform, template matching and active contour deformation variations are proposed in order to perform facial feature detection. A geometric method for computing the gaze direction based on the symmetry properties of certain facial features is also proposed. The proposed methods have provided rather good facial feature detection. Examples demonstrating the success of the various proposed methods as well as criticism when they fail have been included.

References

[1] K. Sobottka and I. Pitas, A novel method for automatic face segmentation, facial feature extraction and tracking, *Signal Processing: Image Communication*, **12**(3), 263-281, (1998).

[2] G. Yang and T.S. Huang, Human face detection in a complex background, *Pattern Recognition*, **27**(1), 53-63, (1994).

[3] C. Kotropoulos and I. Pitas, Rule-based face detection in frontal views, *Proc. ICASSP '97*, 2537-2540, (1997).

[4] R. Chellapa, C.L. Wilson, and S. Sirohey, Human and machine recognition of faces: A survey, *Proceedings of the IEEE*, **83**(5), 705-740, (1995).

[5] R. Brunelli and T. Poggio, Face recognition: Features versus Templates, *IEEE Trans. Pattern Analysis Mach. Intell.*, **15**(10), 1042-1052, (1993).

[6] J. Illingworth and J. Kittler, The Adaptive Hough Transform, *IEEE Trans. Pattern Analysis and Mach. Intell.*, **9**(5), 690-698, (1987).

[7] R.M. Haralick and L.G. Shapiro, *Computer and Robot Vision*, vol. I, Addison-Wesley, Reading, Massachussets, (1992).

[8] X. Li and N. Roeder, Face contour extraction from front-view images, *Pattern Recognition*, **28**(8), 1167-1179, (1995).

[9] S. Fischer and B. Duc, Shape normalization for face recognition, *Lecture Notes on Computer Science: AVBPA '97*, Crans-Montana, Switzerland, 21-26, (1997).

[10] M. Kass, A. Witkin and D. Terzopoulos, Snakes: Active contour models, *International Journal of Computer Vision*, **1**(4), 321-331, (1988).

[11] S.R. Gunn and M.S. Nixon, Snake Head Boundary Extraction Using Global and Local Energy Minimisation, *Proc. ICPR '96*, 581-585, (1996).

[12] C. Nastar and N. Ayache, Fast Segmentation, Tracking, and Analysis of Deformable Objects, *Proc. ICCV '93*, 275-279, (1993).

[13] D.J. Williams and M. Shah, A Fast Algorithm for Active Contours and Curvature Estimation, *CVGIP: Image Understanding*, **55**(1), 14-26, (1992).

[14] K.M. Lam and H. Yan, An Improved Method for Locating and Extracting the Eye in Human Face Images, *Proc. ICPR '96*, 411-415, (1996).

[15] B. Esme, B. Sankur and E. Anarim, Facial Feature Extraction Using Genetic Algorithms, *Proc. ICPR '96*, 1511-1514, (1996).

[16] M.J.T. Reinders, P.J.L. van Beek, B. Sankur and J.C.A. van der Lubbe, Facial feature localization and adaptation of a generic face model for model-based coding, *Signal Processing: Image Communication*, **7**(1), 57-74, (1995).

[17] I. Pitas, *Digital Image Processing Algorithms*, Prentice Hall, UK, (1993).

features	correctly detected [%]
left cheek	78
right cheek	86
chin	65
left eyebrow	73
right eyebrow	81

Table 1: Detection percentage for facial feature extraction

List of Figures

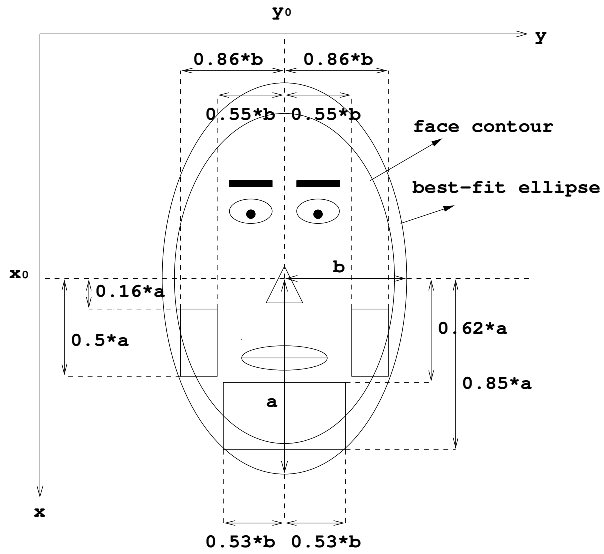

Figure 1: Illustration of relevant subimages for Hough transform.

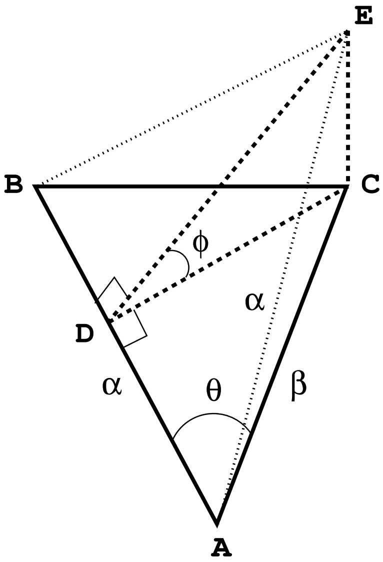

Figure 2: Geometry of projected isosceles triangle.

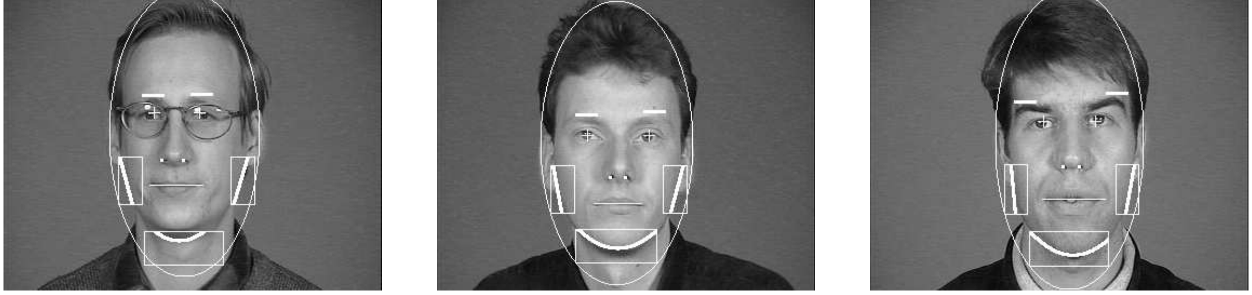

(a) (b) (c)

Figure 3: Examples of good results for facial features.

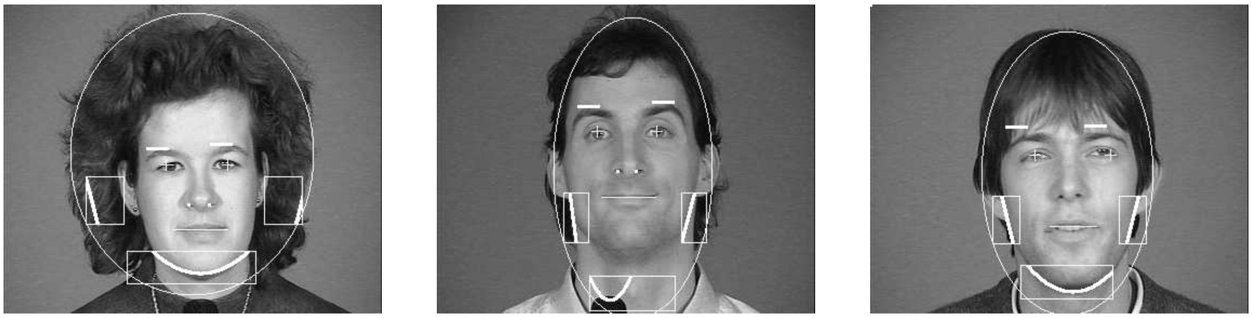

(a) (b) (c)

Figure 4: Examples of erroneous results for facial features.

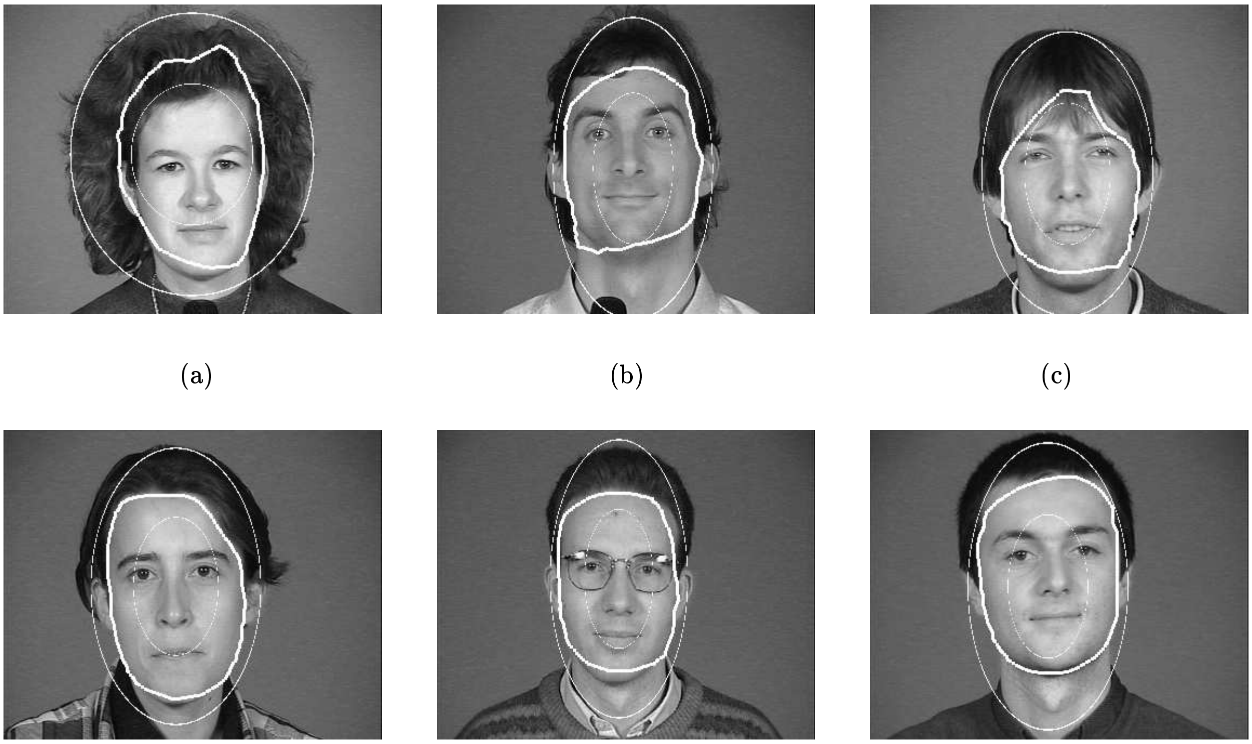

(a) (b) (c)

(d) (e) (f)

Figure 5: Examples of results for face contour.

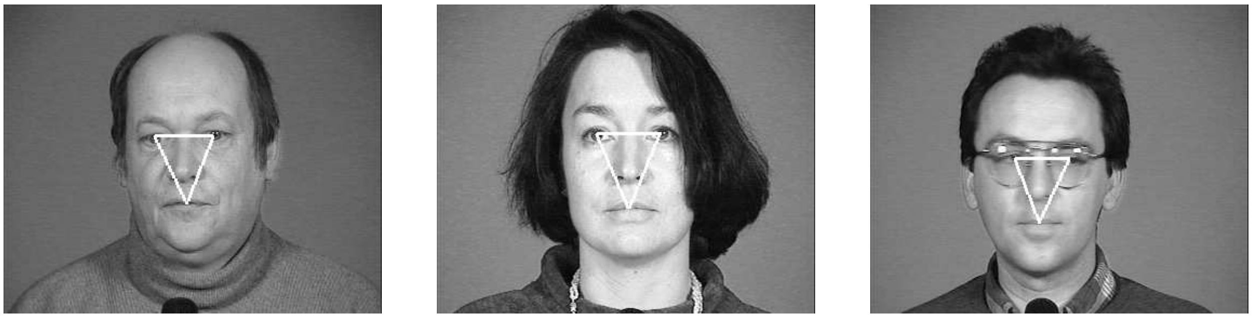

(a) (b) (c)

Figure 6: Examples of results for gaze direction.