

A Face Image Identification Using Artificial Neural Networks

Mr. Sumit ChauhanScholar, M.Tech,
UPTU, Lucknow,
UP, India.**Dr.Yatin Agrawal**Associate Professor,
Dept of CSE,
GNIOT, Gr. Noida, UP, India.**Mr.Neeraj Kumar Verma**Assistant Professor,
Dept of CSE,
GNIOT, Gr. Noida, UP, India.

Abstract:

Face Image Identification is a part of many face identification systems, due to its ability to focus on computational resources to the part of an image having a face. The task of detecting and locating a human faces in a given images is quit complex due to the variations present across human faces, including skin colour, position, expression, pose and orientation, and the presence of 'facial furniture' like glasses or facial hair. Differences in camera angle, lighting conditions and image resolution further complicate the situation. A Review of some of the different approaches to the problem has been included in this report as a preliminary introduction of image identification.

An implementation of one of the identified image-based approaches to the problem using Artificial Neural Networks, has been included, explained and analysed. The parameter choices are investigated to confirm an optimal set of operational parameters. A small adjustments were made to the Image Identification system, which reflecting some identified limitations and leading to small but significant improvements. The rate of a face Image Identification was increased from 29% to 44% on the first test set, and from 44% to 57% on the second test set, at the cost of an increase in the number of false positives. A result analysis of the findings is included as a summary and details of further work.

Keywords:

Eigen vector, facial furniture, ANN, Face train, Face scan.

1. Introduction:

Face identification is a trivial task for the human brain has proved to be extremely difficult to implement artificially. Face Image Identification involves separating image windows into two classes; one containing faces (targets), and other containing the background (clutter). It is difficult because although similarities exist between faces, they can vary considerably in terms of age, skin tone and facial expressions. The problem is further complicated by differing lighting conditions, image qualities and geometries, as well as the possibility of partial occlusion and disguise. An ideal for face detector would be able to detect the presence of any face under any set of lighting conditions and background. For basic pattern recognition systems, although some of these effects can be avoided by assuming and ensuring a uniform background and fixed uniform lighting conditions. These assumptions are acceptable for some applications such as the automated separation of nuts from screws on a production line, where lighting conditions can be controlled, and the image background will be uniform. For many applications however, this is unsuitable, and systems must be designed to accurately classify images subject to a wide variety of unpredictable conditions.

2 Literature review

2.1 Face Image Identification approaches – an Overview

A review is conducted on Face Image Identification techniques, and identified two broad categories named feature-based approach and image-based approach.

2.1.1 Feature-Based Approach

S.Hjelm & Low [1] divide various feature-based systems into three broad sub-categories: low-level analysis, feature analysis, and active shape models.

2.1.1.1 Low-level Analysis

Low-level Analysis works with the division of visual features using the various properties of the pixels, predominantly gray-scale or colour. Edge detection (detecting morphological changes in pixel properties) was first implemented by Sakai et al [2] for detecting features of face in line drawings. Craw et al [3] developed this further to trace outline of a human head, allowing feature analysis to be constrained to within the head outline.

2.1.1.2 Featured Analysis

Feature analysis recognize the face by using facial features for example a pair of dark regions found in the face area increase the probability of a facial existence. The facial feature extraction algorithm [4], is a good example of feature searching, achieving 82% accuracy with invariance to gray and colour information, failing to detect faces with glasses and hair covering the forehead.

2.1.1.3 Active shape models

Active shape models represent the actual physical and hence higher-level appearance of features. These models are released near to a feature, such that they interact with the local image, deforming to take the shape of the feature. There are three types of active shape models that have been used throughout the literature: snakes, deformable templates and smart snakes.

2.2 Image Based Approach

Face Image Identification by explicit modelling of facial features is a very rigid approach which has been shown to be troubled by the unpredictability of faces and environmental conditions.

There is a need for more robust techniques, capable of performing in hostile environments, such as detecting multiple faces with clutter-intensive backgrounds. Hjelm s and Low [1] divide the group of image-based approaches into three sections: Linear subspace methods, Artificial Neural Networks, and statistical approaches.

2.2.1 Linear subspace Method

Images of human faces lie in a subspace of overall image space which can be represented by methods closely related to standard multivariate statistical analysis, including Principal Component Analysis (PCA), Linear Discriminant Analysis (LDA), and Factor Analysis (FA). In the late 1980s, Sirovich and Kirby [6] developed a technique using PCA to represent human faces. The technique first finds the principal components of the distribution of faces (expressed in terms of eigenvectors). Each face in the set can then be approximated by a linear combination of the largest eigenvectors, more commonly referred to as eigenfaces.

2.2.2 Artificial Neural Network Approach

Early approaches based on the simple Multiple Layer Perceptrons (MLP) gave encouraging results on fairly simple datasets. The first advanced neural approach which reported Result statistics on a large, visually complex dataset, was by Rowley et al [7]. Their system incorporates face knowledge in the retinally connected Artificial Neural Network architecture, with specialised window sizes designed to best capture facial information (e.g. horizontal strips to identify the mouth). Images are pre-processed before being classified by the network, the output from which is post-processed to remove overlapping Image Identifications, resulting in one Image Identification per face, and a reduction in false positives.

2.2.3 Statistical Approaches

A system based on support vector machine is an example of Image-Based approaches that do not fit into either of the other categories. In Osuna et al [8] a support vector machine (SVM) is applied to face Image Identification. A SVM with a 2nd degree polynomial as a kernel function is trained with a decomposition algorithm which guarantees global optimality.

Images are pre processed and trained with a bootstrap learning algorithm (more detail in section 2.2.6). Other research into SVMs has attempted to improve the Osuna detector [8]

3. Neural Network-Based Detector Implementation

There are two main functions: ‘facetrain’ to create and train a Artificial Neural Network and ‘facescan’ to scan new images for faces. A description of each function is given below

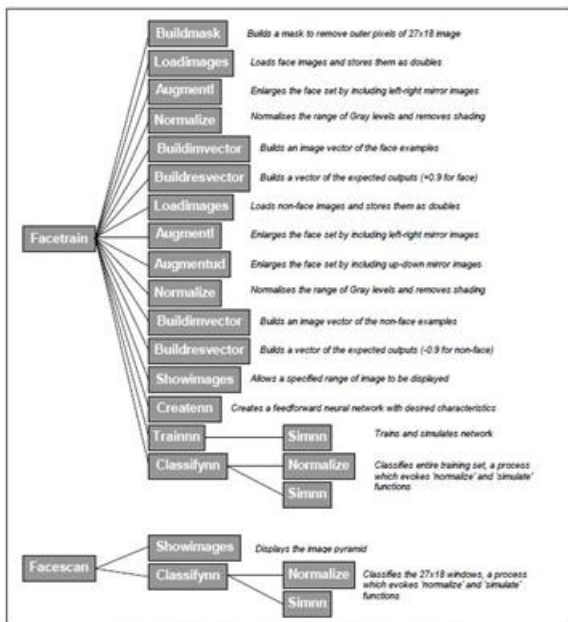


Fig 1. The hierarchy of the Original Sanner detector

Facetrain

A set of 27x18 images from a training set is loaded and stored as an image vector. There are two vectors, one which contains numerous face examples, the other for non-face examples. Each image vector is then augmented, adding mirror-images of the original training examples, to create a larger training set. A mask is applied to the face examples, removing pixels outside of the oval mask to focus the attention of the classifier on the true face region. Pixels in the unmasked area are then normalized.

Facescan

Once the system has been created and trained, it is possible to classify new unseen images. The second function, “facescan”, conducts the final task, scanning previously unseen images for faces. Images are processed prior to classification, which involves the construction of an image resolution pyramid, and scanning 27x18 window regions, normalising each window before passing it to the network for classification. The image resolution pyramid is used to allow faces of differing scales (sizes) to be detected. When calling the ‘facescan’ function, a number of parameters can be specified which control the number of levels in the pyramid and the scale factor for resizing between levels, as well as other parameters specifying the network and mask to be used, and a threshold value, above which images are classified as faces.

4. Experimental Work

This section details work carried out to measure the Result of the discussed Sanner face detector, and to analyse the improvements made.

4.1 Result Analysis- Original Detector

The Result of the original face detector developed by Sanner will be discussed, and a set of optimal values for the various tuning parameters will be investigated. All the experimental work is to be carried out in Matlab using the existing code written by Scott

Sanner. Some additional scripts will be written to implement any improvements, and to automate some of the Result testing experiments which would otherwise be a tedious repetitive procedure.

4.2 Classification result

At the very heart of the system lies the classifier, the object that actually makes the decision as to whether an image is receives contains a face or not. Initially tests were carried out to investigate how well the classifier could classify the data set on which it was trained. Although this is not indicative of true Result, it serves as a guide to how well the network is learning from the Training Data Set. The classify function written by Sanner, classifies 27x18 pixel images as either face or non-face images, producing a numerical value between -0.9 and 0.9 (0.9 strongly indicating the presence of a face, -0.9 indicating the absence of a face).

The addition of a threshold value allows the classifier to be tuned somewhat, such that an image is marked as when the numerical output value exceeds the threshold. A script entitled “massclassify” makes use of the classify function by repeatedly classifying each example from the Training Data Set. For these experiments a virtual threshold of ‘0’ is assumed such that anything classified positive is defined as a face image, and negative values are defined as scenery images.

This additional code is included in the appendices for reference. Face and non-face examples are classified separately. Whilst processing the face examples, correct classification values (exceeding the virtual threshold) are used to increment a ‘face counter’. Likewise a ‘non-face counter’ tallies the number of times non-face examples that are below the threshold value (again correct classification).

Upon completion the classification rate and number of incorrect classifications for both the face and non-face examples are displayed. Figure 2 shows

the results of ‘massclassify’ when used to classify the original Training Data Set set using Sanner’s original detector.

Number of Faces Examples:	30
Correct Classification Rate:	93.3%
Number of Incorrect Classifications:	2
Number of Scenery Examples:	40
Correct Classification Rate:	97.5%
Number of Incorrect Classifications:	1

Figure 2 – Original Face Detector Learning Result Statistics

a. Mask

Literature Review. The unanimous acceptance of this mask throughout various approaches infers that the mask is somewhat optimal already, and thus no experiments will be done The mask is used to remove pixels towards the edge of the 27x18 images, thus focusing the attention of the network on the unmasked oval region, most likely to contain a face. The hosen mask is shown in the appendices and closely mirrors masks chosen by many others in the with other alternatives masks for the purposes of this project.

4.5 Network

Various characteristics of the network can be changed or varied including the network type, number of hidden nodes, training algorithm used and the training duration. Each parameter will be taken in turn and analysed.

4.6 Network Type

Changing the type of network used could potentially improve the Result of the detector, although the chosen feed-forward type is an excellent choice for this type of application, a choice which is mirrored by other Artificial Neural Network based face Image Identification systems including Rowley et al [7].

Therefore due to the widespread acceptance of this network type, making changes at this stage is deemed unnecessary. It is thought that any complexity of problem can be solved with just a single layer of

hidden neurons. With a greater number of hidden neurons, there are more weights to tune during training, and thus a more complex a decision boundary can be formed (although too many neurons can lead to over fitting of the boundary to the training set, thus poor generalization. The number of hidden neurons will be varied from 1 through to 1000 (25 being the default number in the original design), and the 'massclassify' function will be used to see how well the system learns the training set.

5.1 Training Duration

During training, the various weights and biases are updated incrementally upon processing the Training Data Set. Training continues until either the Result function reaches a specified goal, or until the number of iterations reaches a pre-defined maximum value.

No. of iterati	Ferr	Nerr	Terr	Fdetect (%)	Ndetect (%)	Tdetect (%)
1	0.958	0.015	0.36	4.2	98.5	63.6
5	0.917	0.191	0.46	8.3	80.9	54
10	0.4	0.191	0.26	60	80.9	73.1
20	0.225	0.034	0.10	77.5	95.6	89.5
50	0.1	0.039	0.06	90	96.1	93.8
75	0.058	0.034	0.04	94.2	96.6	95.7
100	0.042	0.029	0.03	95.8	97.1	96.6
200	0.033	0.02	0.02	96.7	98	97.5
300	0.017	0.005	0.00	98.3	99.5	99.1
400	0	0.005	0.00	100	99.5	99.7
500	0.008	0	0.00	99.2	100	99.7
1000	0	0	0	100	100	100

Figure 3 – How varying the maximum number of iterations affects ability to learn training set.

This maximum value was set as 500 for Sanner's detector [9], but has been provided as a parameter to the 'trainnn' function to allow this value to be adjusted. The maximum number of iterations will be varied from 1 to 1000, retraining the network for each value, and monitoring the Result statistics. The results are presented in above Figure. Although Figure 3 shows the training error falls to zero with 1000 iterations, this merely reflects that the network can perfectly classify the set on which it was trained, not a true indicator of real Result. As Figure 3 and 3b show, the recommended number of iterations, 500, is a good choice. Another parameter, the concept of which has been introduced previously, is a parameter that defines

the numerical level above which an image must 'score' to be deemed a 'face'. The threshold parameter is passed to the face scan function. When a new, previously unseen image is to be scanned with the function, a resolution. pyramid is constructed, and each level is divided into 27x18 pixel windows. . Each window is then classified individually, and those which produce a numerical value

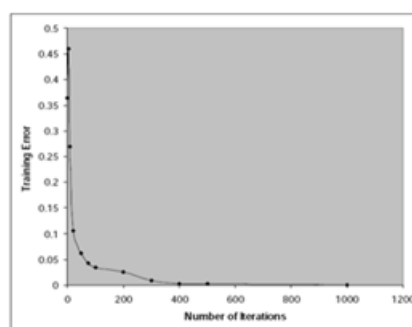


Figure 3b – The affect of increasing the number of iterations on training error.

greater than the threshold are declared to be faces, and are outlined with a black bounding box indicating the location of the face. The threshold value therefore defines the selectivity of 'facescan' function. Sanner [9] suggests that a threshold value between 0.4 and 0.6 is optimal, a statement that should be confirmed by these findings.

		Thr 0.0	Thr 0.1	Thr 0.2	Thr 0.3	Thr 0.4	Thr 0.5	Thr 0.6	Thr 0.7	Thr 0.8	Thr 0.9
Pos. Rate											
PP1	1	1	1	1	1	1	1	1	1	1	1
PP2	1	1	1	1	1	1	1	1	1	1	1
PP3	1	1	1	1	1	1	1	1	1	1	1
PP4	1	1	1	1	1	1	1	1	1	1	1
PP5	1	1	1	1	1	1	1	1	1	1	1
PP6	1	1	1	1	1	1	1	1	1	1	1
PP7	1	1	1	1	1	1	1	1	1	1	1
PP8	1	1	1	1	1	1	1	1	1	1	1
PP9	1	1	1	1	1	1	1	1	1	1	1
PP10	1	1	1	1	1	1	1	1	1	1	1
NP1	0	0	0	0	0	0	0	0	0	0	0
NP2	0	0	0	0	0	0	0	0	0	0	0
NP3	0	0	0	0	0	0	0	0	0	0	0
NP4	0	0	0	0	0	0	0	0	0	0	0
NP5	0	0	0	0	0	0	0	0	0	0	0

Figure 4 – A table of thresholding experiments with the original Sanner detector.

The results are tabulated in Figure 4, showing how Result (Image Identification rate relative to number of false positives) varies as the threshold is changed from 0 to 0.9. The following results were obtained by

repeatedly scanning a test set containing 13 images – 8 face/5 non-face, taken from a dataset compiled by Rowley et al [7], at CMU [10]. A script, ‘threshold’, was written to automate this process. The results from Figure 4 are presented in graphical form in Figure 4b. The plot highlights the trade off between correct Image Identification rate and the number of false positives. From figures 4 and 4b a threshold of 0.7 has been found to be optimal, that is to say best Image Identification rate relative to number of false positives.

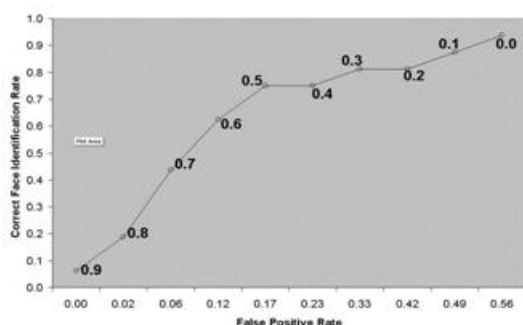


Figure 4b – An ROC Plot for the Original Face Detector.

5.2 Pyramid

When images are scanned for faces (using the facescan function), an image pyramid is constructed. This pyramid, designed to facilitate the Image Identification of faces at different scales, is formed of a discrete number of levels. The original image is resized by a pre-defined scale factor for each level of the pyramid. Sanner provided for the adjustment of three important parameters regarding the construction of the pyramid, each of which will be analysed in turn.

5.3 No. of Pyramidal Levels

The greater the number of levels, the more faces at distance from the camera can be detected, at the expense of an increase in the number of false Image Identifications.

With fewer levels, the scale of possible face Image Identifications will be reduced, but similarly the number of false positives is likely to decrease. Sanner [9] suggests that 6 levels for the pyramid is optimal,

an assertion mirrored by others working in the field (see section 2). To confirm this optimal value, two images have been created containing faces of different sizes. Both the images were constructed by combining a number of images collected by Rowley et al for their face detector [10]. Each of the images was scanned using the facescan function, with the “Levels” parameter varied between 1 and 20. The results are shown in figure 5.

No. Levels	SCALE01 (6 faces)		SCALE02 (31 faces)	
	F d	Num fd	Fdet	Num
1	0	9	1	18
2	2	16	8	26
4	4	30	18	33
6	4	37	20	48
8	4	37	22	51
10	4	37	22	54
12	4	39	27	67
14	4	40	27	69
16	4	42	27	73
18	4	45	27	76
20	4	49	27	81

Figure 5 – The affect that varying the number of pyramid levels has on classification Result (face Image Identification rate Fdetect, no. false positives Numfd)

The original image is resized to form the various layers of the resolution pyramid. Each level is resized by a specified factor of the previous level, thus this ‘scale factor’ also influences the range of scales of faces that can be detected. Sanner [9] suggested an optimal value of 1.2 what appears to be a popular choice throughout the literature. The following tests are expected to agree with this value. The same two images were scanned using once more, this time varying the scale factor from 1.1 to 1.5 at 0.1 increments. All other parameters were kept constant. A face Image Identification rate of 93.3% with only one misclassified scenery image was reported. These statistics only indicate how well the network is able to classify the Training Data Set, and thus how well it has “learned” the training set during training. An insight into the true Result of the system can only be achieved through the use of an independent test set of

“unseen” images, producing statistics more indicative of system Result in real world applications. No standard test set or test procedure exists, and thus comparison between this system, and other proposed face Image Identification systems is almost impossible, an observation made in the Literature Review. In order to produce a realistic reflection of Result in the real world, a test set should evaluate the Result of the detector under a wide range of conditions, representing the true nature of the variation fraught environment. The test set should contain images of various resolutions, with both face and scenery examples present.

Start Level	SCALE01 (6 faces)		SCALE02 (31 faces)	
	F _{det}	Num	F _{detec}	Num _{fd}
1	4	40	23	41
2	4	32	23	34
3	3	24	22	32
4	3	17	20	25
5	2	12	6	20
6	0	7	4	14
7	0	4	4	12
8	0	3	4	9
9	0	1	4	8
1	0	0	1	4

Figure 6 – The affect of varying the start level on classification Result.

Faces should be present at numerous scales, under various lighting conditions, and with differing complexities of background. The majority of faces should be frontal faces (as this detector is designed to detect frontal faces), although slight rotations can be introduced. Two test sets have been chosen to evaluate the Result of Sanner's Face detector [9]. The first, containing 42 face images and 0 scenery images, is thought to be easier than the second test set containing 52 face images, and 13 fairly complex scenery images. Both test sets were obtained from CMU [10], and contain images used in the evaluation of the Rowley et al [7] detector (a summary of the images in the two test sets is included in section 9).

Two scripts ('testset01' and 'testset02') have been written to automate the scanning of the training images. Each image is scanned individually using a pyramid (levels = 6, scale-factor = 1.2). Bounding

boxes are drawn around image windows which exceed the threshold (0.8) and a resultant image is returned and stored for manual inspection. Figure 6 shows the Result of Sanner's original detector [9] when analysed with each of the two test sets.

Start Level	SCALE01 (6 faces)		SCALE02 (31 faces)	
	F _{det}	Num	F _{detec}	Num _{fd}
1	4	40	23	41
2	4	32	23	34
3	3	24	22	32
4	3	17	20	25
5	2	12	6	20
6	0	7	4	14
7	0	4	4	12
8	0	3	4	9
9	0	1	4	8
1	0	0	1	4

The number of images in these test sets is ambiguous; it depends on whether you count partially occluded faces, profile faces, illustrations of faces, and animal faces. For the purposes of this project, 'testset01' contains 178, and 'testset02' contains 180. The results show that the original detector detected 29% of the faces in 'testset01' with an average of 11 false positives.

Test Set 01			
Image Number	No. Faces In Image	No. Faces Detected	No. False Positives
1	5	5	100
2	1	0	18
3	4	0	100
4	2	0	100
5	1	1	100
6	2	1	50
7	2	0	24
8	1	0	26
9	1	1	50
10	1	1	32
11	1	0	24
12	1	1	30
13	1	1	16
14	1	1	15
15	1	0	26
16	25	22	100
17	1	1	11
18	4	0	200
19	5	4	100
20	1	0	26
21	1	0	35
22	2	0	36
23	1	0	35
24	1	0	12
25	1	0	4
26	2	0	3
27	1	0	13
28	2	2	16
29	2	0	30
30	2	0	30
31	6	5	30
32	6	4	30
33	6	2	14
34	7	0	15
35	6	1	100
36	12	6	75
37	4	3	150
38	14	2	50
39	1	1	40
40	1	1	40
41	9	1	50
42	9	5	35

Figure 7 –Face Detector Result Evaluation with replaced training set
Result Analysis for Detector Trained with Replaced Training Set

Test Set 01

- Face Image Identification Rate: 44%

- Average Number of False Positives: 18

Test Set 02

- Face Image Identification Rate: 57%
- Average Number of False Positives: 46

Test Set 02			
Image Number	No. Faces In Image	No. Faces Detected	No. False Positives
1	2	2	50
2	1	0	200
3	1	0	35
4	1	0	40
5	1	1	30
6	1	0	25
7	1	0	30
8	6	3	30
9	1	0	1
10	1	1	4
11	0	0	15
12	1	1	2
13	57	42	50
14	0	0	24
15	1	0	25
16	1	0	60
17	1	1	20
18	6	5	20
19	2	2	22
20	6	2	23
21	0	0	8
22	0	0	60
23	1	0	29
24	2	2	10
25	0	0	60
26	1	0	17
27	1	0	20
28	1	0	2
29	1	0	27
30	2	1	300
31	9	2	51
32	4	3	60
33	1	0	20
34	0	0	60
35	2	2	60
36	3	0	42
37	1	0	300
38	1	0	42
39	1	0	63
40	14	11	23
41	8	2	60
42	2	2	120
43	1	0	17
44	0	0	24
45	3	3	60
46	2	1	60
47	1	0	45
48	1	0	40
49	4	4	28
50	1	1	15
51	0	0	27
52	1	1	29
53	1	0	6
54	1	1	6
55	0	0	120
56	1	0	25
57	1	0	25
58	0	0	60
59	0	0	53
60	0	0	9
61	2	1	32
62	2	1	28
63	1	0	8
64	0	0	30
65	1	0	2

NB: Numbers in bold (50) indicate an estimate of the number of Image Identifications. As the images are inspected manually, it with large numbers of overlapping Image Identifications, an accurate number is impossible to determine and thus an estimate is

suggested. There were two attempts to improve the system Result: the first entirely replaced the Training Data Set set, the second complemented the new set with additional image with the aim of improving the Image Identification rate slightly, but mainly to reduce the number of false positives. The network was retrained with the new image datasets, and then used to scan the same two test sets as the original Image Identification system.

	Original Detector	Replaced Dataset	Improved Dataset
Testset01 – Fdetect	2	43%	43%
Testset01 – Numfd	1	47	42
Testset02 – Fdetect	4	57%	58%
Testset02 – Numfd	1	46	44

Figure 8 – Comparison of Result statistics

The results (figure 8 and 8b) were encouraging with a substantial increase in Image Identification rates, although an accompanying increase in the number of false Image Identifications. In the first instance, the Training Data Set was completely replaced resulting in many more faces being detected. The number of false positives indicated a clear problem with the Training Data Set, and thus improvements were made to the training set. Despite these attempts, the improvement was fairly small, as the number of additional images added to the training set was small in comparison to the size of the set, and thus contributed little to the overall “location” of the class decision boundary. This highlights the importance of the non-face examples, the difficulty in constructing a representative set, and the need for a better technique for doing so. Face Image Identification has a many applications including Security applications, Until recently, much of the work in the field of computer vision has focussed on face identification, with very little research into face Image Identification.

Human face Image Identification is often the first-step in the identification process as detecting the location of a face in an image, prior to attempting identification can focus computational resources on the face area of the image. Although a trivial task that human perform

effortlessly, the task of face Image Identification is a complex problem in computer vision, due the great multitude of variation present across faces.

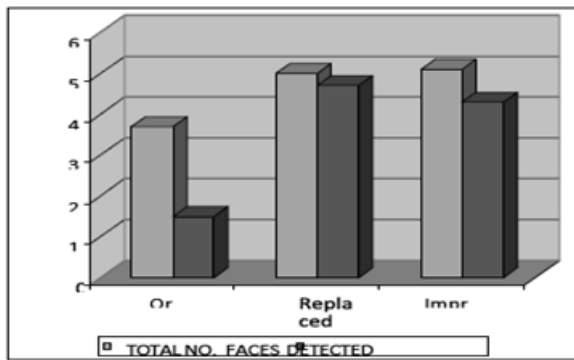


Figure 8b – A graph comparing the performance.

In conclusion, the system developed by Sanner [5] is a good example of a Artificial Neural Network based system, indicative of some of the more complex detectors in the field. It mirrors the strengths of the technology providing impressive classification results from a relatively small image training set, and also reflects the major limitations, mainly computational expense, and reliance on the Training Data Set. It illustrates well some of the key problems that developers of intelligent artificial face Image Identification systems are faced with, not only in the field of Artificial Neural Networks but across the board.

Reference

- 1.Hjelm s, E. and Low, B.(2001). Face Image Identification: A Review. Computer Vision and Image Understanding. 83,236-274.
- 2.Sakai, T., Nagao, M. and Kanade, T. (1972). Computer analysis and classification of photographs of human face. Proc. First USA – Japan Computer Conference,2.7.
- 3.Craw, I., Ellis, H. and Lishman J.R.(1987). Automatic extraction of face- feature. Pattern Recog. Lett.183-187.
- 4.De Silva, L.C., Aizawa, K., and Hatori, M.(1995). Image Identification and tracking of facial features by using a facial feature model and deformable circular

template. IEICE Trans. Inform Systems. E78-D(9),1995-1207.

5.Huang, C.L., and Chen, C.W.(1992). Human facial feature extraction for face interpretation and identification. Pattern Recog, 25,1435-1444.

6.Sirovich, L. and Kirby, M. (1987). Low-dimensional procedure for the characterization of human faces. J. Opt. Soc. Amer., 4,519-524.

7.Rowley, H.A., Baluja, S., and Kanade, T.(1998). Artificial Neural Network-based face Image Identification. IEEE Trans. Pattern Anal. Mach. Intelligence, 20,23-38.

8.Osuna, E., Freund, R., and Girosi, F. (1997). Training support vector machines: An application to face Image Identification. IEEE Proc. of Int. Conf. on Computers Vision and Pattern Recog, 33, 1369-1382.

9.Sanner, S. "Rowley-Baluja-Kanade Face Detector". Available online at <http://www.cs.toronto.edu/~ssanner/Software/Vision/Project.html>

10.Rowley, H.A., Baluja, S., and Kanade, T.(1998). CMU Image Dataset. Available online at <http://www2.cs.cmu.edu/~har/faces.html> <http://www.fuzzgun.btinternet.co.uk/or> <http://sourceforge.net/projects/rodney/>