UNIVERSITY OF MINNESOTA

CENTER FOR TRANSPORTATION STUDIES

INTELLIGENT TRANSPORTATION SYSTEMS INSTITUTE

Detecting Driver Fatigue Through the Use of Advanced Face Monitoring Techniques

Final Report

Prepared by

Harini Veeraraghavan
Nikolaos Papanikolopoulos

Department of Computer Science and Engineering
University of Minnesota

CTS 01-05

HUMAN CENTERED TECHNOLOGY TO ENHANCE SAFETY AND TECHNOLOGY

Technical Report Documentation Page

1. Report No. CTS 01-05	2.	3. Recipient's Accession No.
4. Title and Subtitle Detecting Driver Fatigue Through the Use of Advanced Face Monitoring Techniques		5. Report Date September 2001 6.
7. Author(s) Harini Veeraraghavan Nikolaos P. Papanikolopoulos		8. Performing Organization Report No.
9. Performing Organization Name and Address Department of Computer Science and Engineering University of Minnesota 4-192 EE/CSci Building 200 Union Street SE Minneapolis, MN 55455		10. Project/Task/Work Unit No. 11. Contract (C) or Grant (G) No. (C) (G)
12. Sponsoring Organization Name and Address ITS Institute Center for Transportation Studies University of Minnesota 200 Transportation and Safety Building 511 Washington Ave. SE Minneapolis, MN 55455		13. Type of Report and Period Covered Final Report 14. Sponsoring Agency Code
15. Supplementary Notes		

16. Abstract (Limit: 200 words)

Driver fatigue is an important factor in many vehicular accidents. Reducing the number of fatigue-related accidents would save society a significant amount financially, in addition to reducing personal suffering. The researchers developed a driver fatigue monitoring system that uses a camera (or cameras) to detect indications of driver fatigue. The mechanism detects and tracks the eyes of the driver based on human skin color properties, along with templates that monitor how long the eyes are open or closed. Tests of the approach were run on 20 human subjects in a simulated environment (the driving simulator at the Human Factors Research Laboratory) in order to find its potential and its limitations. This report describes the findings from these experiments.

17. Document Analysis/Descriptors driver fatigue detection, face recognition, vision-based tracking, human factors, driver simulator experiments		18. Availability Statement No restrictions. Document available from: National Technical Information Services, Springfield, Virginia 22161	
19. Security Class (this report) Unclassified	20. Security Class (this page) Unclassified	21. No. of Pages 31	22. Price

Detecting Driver Fatigue Through the Use of Advanced Face Monitoring Techniques

Prepared by:

Harini Veeraraghavan
Nikolaos P. Papanikolopoulos
Artificial Intelligence, Robotics, and Vision Laboratory
Department of Computer Science and Engineering
University of Minnesota
Minneapolis, MN 55455

Published by:

ITS Institute
Center for Transportation Studies
200 Transportation and Safety Building
511 Washington Ave. SE
University of Minnesota

September 2001

CTS 01-05

TABLE OF CONTENTS

LIST OF FIGURES & TABLES

EXECUTIVE SUMMARY

Driver fatigue resulting from sleep deprivation or sleep disorders is an important factor in the increasing number of accidents on today's roads. The purpose of this project was to advance a system to detect fatigue symptoms in drivers and produce timely warnings that could prevent accidents.

This report presents an approach for real-time detection of driver fatigue. The system consists of a video camera directly pointed towards the driver's face. The input to the system is a continuous stream of images from the video camera. The system monitors the driver's eyes to detect micro-sleeps (short periods of sleep lasting 3 to 4 seconds). The system can analyze the eyes in each image as well as compare two frames.

The system uses skin color information and blob analysis for detecting the face pixels of the driver. The system extracts the eye templates of the driver first by taking a difference of two frames and performing blob operation on the difference image. A pattern-matching technique is then used for detecting whether the eyes of the driver are open or closed. If the eyes remain closed for a certain period of time (3 to 4 seconds), the system determines that the person has fatigue and gives a warning signal. The system also checks for tracking errors; once an error is detected, the system returns to the face detection stage.

The system uses a Pentium Pro 200 MHz personal computer with a Matrox Genesis imaging board that holds a Texas Instruments TMS320C80 DSP chip. The performance of the system is two frames per second for tracking and fatigue detection.

The system was tested on 20 different human subjects with different skin color, facial hair, and gender. The system gave very good results for all of the cases except one. The system rarely lost track of eyes for relatively low head displacements, while also detecting the blinks of the driver and giving no false alarms in 19 of the 20 cases studied. The failure in the one case was due mainly to the lighting conditions, as discussed in

detail in the report. The skin color of the individuals affected the performance of the system slightly. For people with fair skin, the system performed very well, as their skin reflects more ambient light.

Chapter 1. INTRODUCTION

A large number of road accidents are caused by driver fatigue. Sleep deprivation and sleep disorders are becoming common problems among car drivers these days. A significant portion of the accidents occurring on U.S. highways is due to driver fatigue. A system that can detect oncoming driver fatigue and issue timely warnings could help prevent many accidents, and consequently save money and reduce personal suffering.

There are many indicators of oncoming fatigue, some of which can be detected using a video camera. One of the symptoms that we try to detect is the “micro-sleep.” Micro-sleeps are short periods in which the driver loses consciousness.

The input to the system are images from a video camera mounted in front of the car, which then analyzes each frame to detect the face region. The face is detected by searching for skin color-like pixels in the image. Then a “blob” separation performed on the grayscale image helps obtain just the face region. In the eye-tracking phase, the face region obtained from the previous stage is searched for localizing the eyes using a pattern-matching method. Templates, obtained by subtracting two frames and performing a blob analysis on the difference grayscale image, are used for localizing the driver’s eyes.

The eyes are then analyzed to detect if they are open or closed. If the eyes remain closed continuously for more than a certain number of frames, the system decides that the eyes are closed and gives a fatigue alert. It also checks continuously for tracking errors. After detecting errors in tracking, the system starts all over again from face detection.

The main focus is on the detection of micro-sleep symptoms. This is achieved by monitoring the eyes of the driver throughout the entire video sequence. The three phases involved in order to achieve this are the following:

1) localization of the face,
2) tracking of eyes in each frame, and
3) detection of failure of tracking.

To make sure that it is the face region that is detected, and not the background that might have skin-like color, the system checks if the area detected as skin has a minimum area (number of pixels). During tracking, the eye templates are matched with the face region to locate eyes. The match scores for the eyes detected are checked continuously. If the match scores for both the open and closed eyes fall below a certain threshold, the system decides that there is an error in tracking and goes back to face detection again. The search space need not be adjusted unless there is a tracking error. The eyes can be detected with fair accuracy unless there is large head-bouncing movement. Further, to ensure correctness in the detection of the eyes, the information about the horizontal alignment and the minimum distance between the two eyes is used.

This report explains the system, tests, and results. Chapter 2 gives an overview of the system; Chapter 3 explains face detection; Chapter 4, eye tracking; Chapter 5 gives the experimental set-up; Chapter 6 gives results; and Chapter 7 closes the report with conclusions and recommendations.

Chapter 2. OVERVIEW OF THE SYSTEM

Functional Description

The system consists of three well-defined phases, namely the face detection, eye tracking and fatigue detection.

The sequences of images from the camera are fed to the system. Initially, the system doesn't know the initial position of the face. The system grabs the first image and tries to find the face region in the image using the skin color model. Due to unfavorable lighting conditions or initial head orientation of the driver, the localization might fail. So the system grabs another frame and repeats the same process until the face region is detected with certainty.

It is assumed that the person's head is not displaced a lot from the previous position between two consecutive frames. So after the face is detected, the eyes are tracked in the region and monitored to detect micro-sleeps. The eye templates obtained previously for the driver are used for localizing the eyes. The eyes are analyzed to determine whether they are open or closed. This information obtained for each frame is passed on to the fatigue detection phase if there is no tracking error. The system doesn't relocalize the eyes unless there is a tracking error. If there is, the area searched is increased in size. So even if the face is displaced significantly between frames, the eyes can still be localized.

The flowchart (Figure 1) gives an outline of the system.

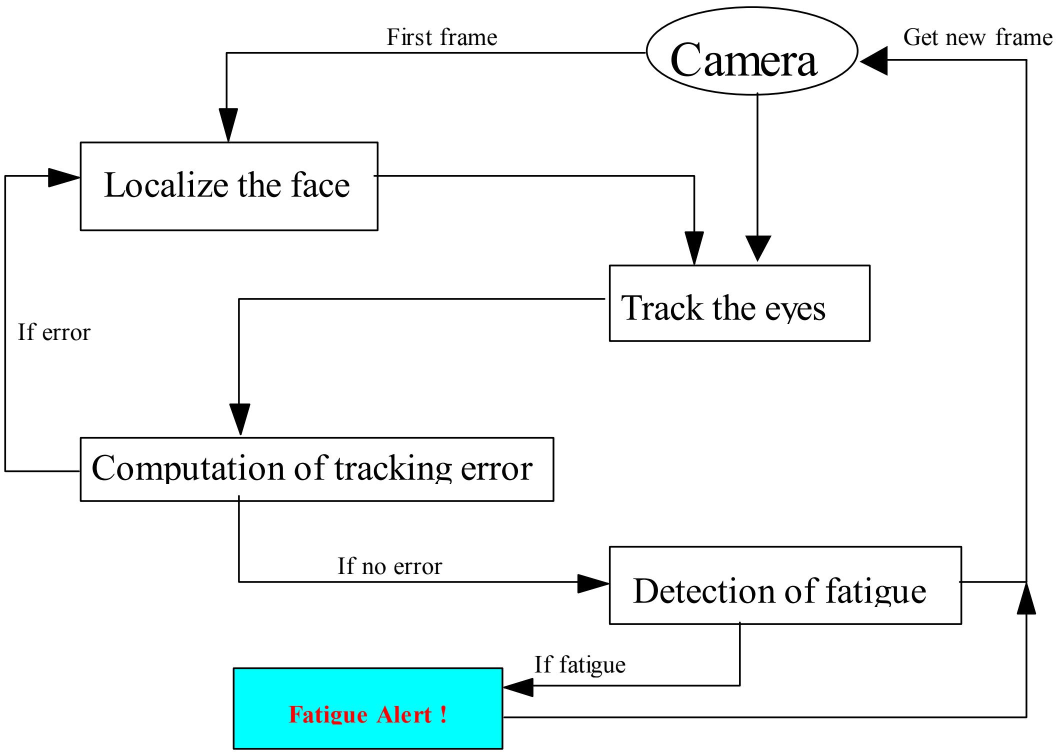

Figure 1. *Flowchart of the approach.*

Once the search region is localized, the face is not tracked again until there is a tracking failure. In the event of failure in tracking, the system goes back to the face localization stage.

Experimental Setup

The system was tested with 20 different subjects. The performance of the system was tested offline using the images from the video camera.

The images were gathered initially from a car simulator (Human Factors Research Laboratory at the University of Minnesota). The set up consisted of a stationary car, which was driven in a simulated environment. The room was darkened except for the light from the camera mounted on the hood of the car and the light reflected from the simulator screen on the face of the driver. This was the only illumination used for obtaining the images. In order for the light from the camera not to bother the driver, it was dimmed. A Panasonic camcorder was used for obtaining the images.

The system for detecting the fatigue uses a Pentium Pro 200 MHz personal computer with a Matrox Genesis imaging board that holds a Texas Instruments TMS320C80 DSP chip.

Chapter 3. FACE DETECTION

Automatic localization or detection of human faces is not easy, as there are difficulties due to the varying lighting conditions, head orientation, and partial occlusions of the face region. We try to detect the face region as a whole unit and use it to extract the information about the position of the eyes in the face.

The face detection was done using the method reported in [19]. The skin color model was used for detecting the face region. Using skin color information to detect the face region is faster and more reliable under constant lighting conditions, as the sizes, head orientation and partial occlusions in the face do not affect the color. This model is useful for detecting the faces of people irrespective of their race and skin color. The skin color though appears to vary over a wide range among people of different races, varies less in color than in brightness. Thus, the skin color model is unaffected by the skin color, eyes, or the facial hair.

In other words, for a RGB representation of an image, the following relation holds for two pixels P_1 with value $[r_1, g_1, b_1]$ and P_2 $[r_2, g_2, b_2]$

$$\frac{r_1}{r_2} \approx \frac{g_1}{g_2} \approx \frac{b_1}{b_2}.$$

The skin pixels have similar color but possibly different brightness. A normalized chromatic color space representation could be used for representing the skin color, as brightness is not important to represent human skin under normal lighting.

Normalized chromatic color representations are defined as the normalized r- and g-components of the RGB color space. This representation removes the brightness information from the RGB signal while preserving its color. Further, the complexity of the RGB color space is simplified by the dimensional reduction to a simple RG color space. The following transformation is used to transforming the input image from RGB space to chromatic color space (r, g)

$$r = \frac{R}{R+G+B}$$

$$g = \frac{R}{R+G+B}$$

In order to distinguish the skin color of the user's face from the other image regions in the image, the distribution of the skin color in the chromatic color space must be known prior to employing the system for detecting the human face. Skin color models vary with the skin color of the people, video cameras used and also with the lighting conditions. Skin pixels clustered in the chromatic color can be represented in chromatic color space using a Gaussian distribution. A skin color model similar to the one developed in [19] was used as the lighting conditions and the camera used were similar.

The skin color distribution using the Gaussian model $N(m, \Sigma^2)$ is represented as

$$\bar{r} = \frac{1}{N}\sum_{i=1}^{N} r_i,$$

$$\bar{g} = \frac{1}{N}\sum_{i=1}^{N} g_i,$$

where, $m = (\bar{r}, \bar{g})$ and

$$\Sigma = \begin{bmatrix} \sigma_{rr} & \sigma_{rg} \\ \sigma_{gr} & \sigma_{gg} \end{bmatrix}$$

Using the skin color model we filter out the incoming video frames to allow only those pixels with high likelihood of being skin pixels. We use a threshold to filter out the skin like pixels from the rest of the image. The filtered image is then binarized and blob operation performed to detect the face region from the rest of the image space. In order to reduce the computational cost and speed up the processing, each incoming frame is sub-sampled to a 160x120 frame.

Chapter 4. EYE TRACKING

The system employs a template matching technique to match the reference eye templates recovered from the previous image with the image area within the estimated face region.

Initialization

The reference eye patterns for each user are recovered previously by taking the difference of two images. The eye blink is used to estimate the position of the eye. The eye templates are recovered by taking a difference of the two images and employing blob (area) operations to isolate the eye regions. For the correct detection of the eye templates, it is required that there is no other motion of the face other than the eye blinks. The eye pattern consists of the eyes centered at the center of the iris of the user.

Using the reference eye patterns, the image is searched for localizing the eyes during the eye-tracking phase.

Eye Tracking

Using the estimated face position detected in the previous frame, the subsequent images are searched for the eyes using the reference templates. For this operation, a grayscale correlation pattern matching method is used. The templates consist of four different images, namely the left open eye, left closed eye, right open eye and right closed eye.

Sample templates are shown in the Figures 2 and 3 below:

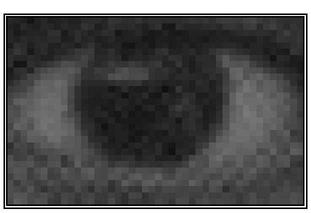

Figure 2. *Open eye template*.

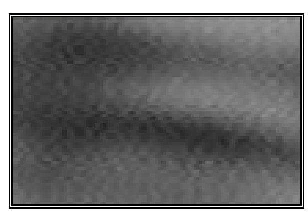

Figure 3. *Closed eye template.*

Eye templates at different head orientations or rotations were not used as this could increase the computational time for each image.

A grayscale correlation method is used. The templates are matched with every pixel in the region being searched and a match score is assigned to each pixel in the target image.

The match ***r*** is computed as:

$$r = \frac{N\sum_{i=1}^{N} I_i M_i - \sum_{i=1}^{N} I_i \sum_{i=1}^{N} M_i}{\sqrt{\left[N\sum_{i=1}^{N} I_i^2 - (\sum_{i=1}^{N} I)^2\right]\left[N\sum_{i=1}^{N} M_i^2 - (\sum_{i=1}^{N} M)^2\right]}}$$

where **N** is the number of pixels in the model, **M** is the model and **I** is the image against which the pattern is being compared.

A match computed by the above expression is unaffected by linear changes in the image or model pixel values. In other words, the search is just as efficient even if the image gets brighter or darker. The value of ***r*** reaches a maximum of **1** when the model matches perfectly with the image or **0** when there is no match.

The actual match scores are computed as:

$$\textit{Score} = \max(r, 0)^2 \times 100\%$$

Acceptance level is assigned for all the matches. Acceptance level is the match score above which a "match" is considered to be found. In other words, if the match score of a pixel with a model being searched is above acceptance level, the model ***M*** is considered to match with the pixel. Otherwise, the match is considered false and we move on to the next model. If a match occurs, we move on to grab the next frame and repeat the process.

The system searches for the open eyes starting from the left eyes first and then looks for the right eyes. If the scores for the open eyes are reasonably higher than the acceptance level and the system decides that the eyes are open, it doesn't search for the closed eye patterns in the image.

In order to avoid mistracking or mismatch, the geometric position of the eyes detected is verified by checking for the horizontal alignment between a pair of eyes as well as their relative distance from each other.

The threshold scores fixed for the open eyes consist of a minimum above which the system decides that the eyes are probably open. When the scores are above the maximum threshold, the system decides for sure that the eyes are open and doesn't search for the closed eyes in the image. But if the scores are between the minimum and the maximum limits, then the system searches the image for closed eyes too in order to remove any mismatches.

The system simultaneously checks for fatigue detection. In case the eyes of the subject remain closed for unusually long periods of time, the system gives a fatigue alert. The fatigue alert persists as long as the person doesn't open his eyes. In case all the matches fail, the system decides that there is a "tracking failure" and switches back to the face localization stage. As the face of the driver doesn't move a lot between frames, we can use the same region for searching the eyes in the next frame.

The following are some snapshots of the performance of the system showing detection of open and closed eyes. The scores for open and closed eyes are at the bottom of the image. The snapshots also show the fatigue alert issued by the system.

Figure 4. *Snapshots of the face of a human subject.*

Figure 5. *Snapshots of the face of a human subject.*

Figure 6. *Snapshots of the face of a human subject.*

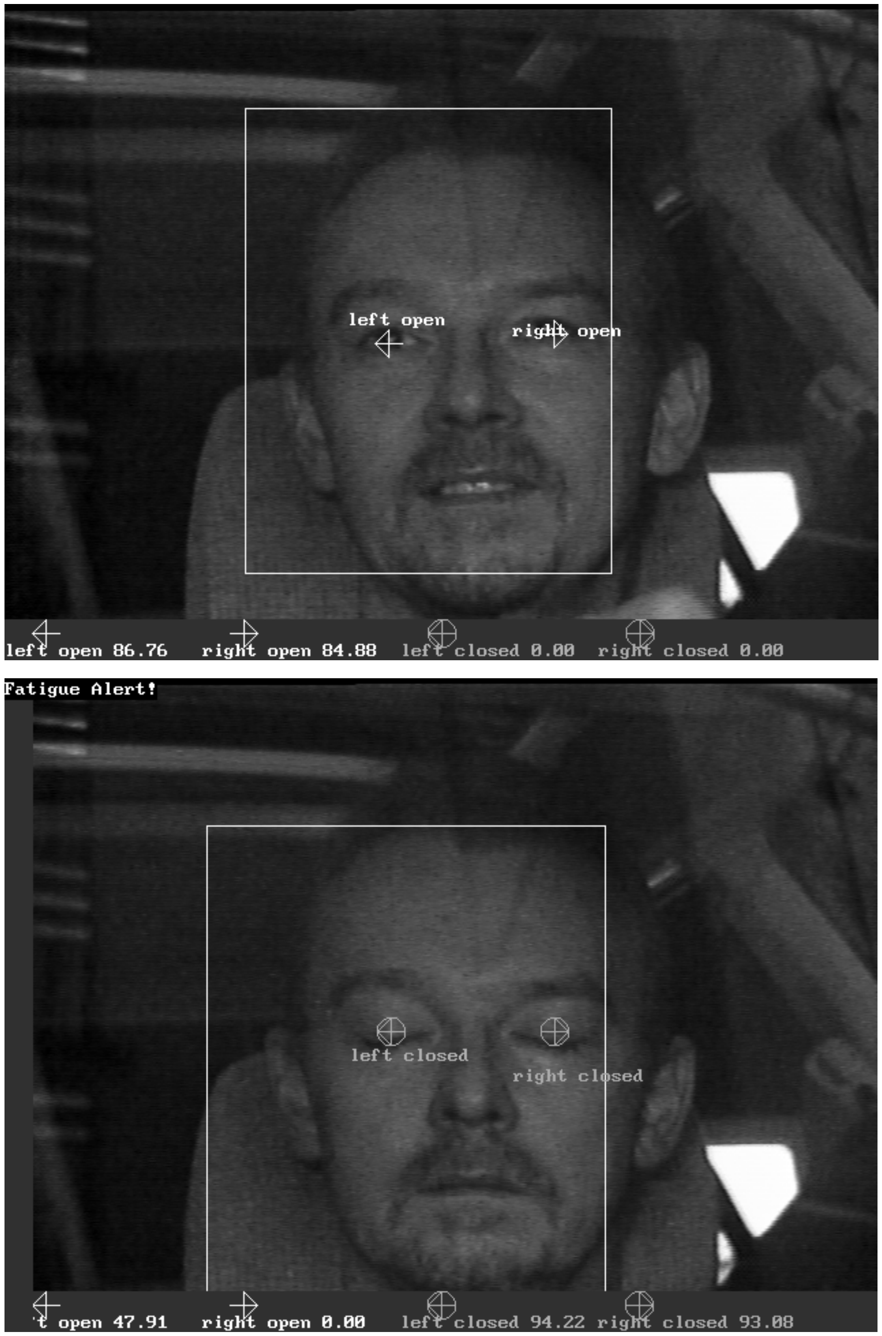

Figure 7. *Snapshots of the face of a human subject.*

Chapter 5. EXPERIMENTAL SETUP

The experiments were carried out in two phases. In the first phase, the images were grabbed using a video camera and in the next phase, these images were used as inputs to the system to detect driver fatigue.

The experiments for grabbing the image were performed in the Human Factors Research Laboratory (University of Minnesota). A flat screen simulator was used for this purpose. The entire experiment for each driver was carried out for 30–40 minutes. The subjects consisted of people belonging to different races, gender and different facial hair. All the subjects considered didn't wear eyeglasses. An appropriate human subject process was established in order to satisfy federal rules.

The simulator consisted of a semi-dark room. The light source used was the light from the camera. The camera used for the test was a Panasonic camcorder. It was mounted on the hood of the car pointing towards the face of the driver. The light from the camera was dimmed so that it didn't cause any discomfort to the driver.

It was found that the light from the camera wasn't sufficient. So, an additional dim light source was used which was placed inside the car.

The simulation used was a graphical simulation of a stretch of road, which repeated itself recursively after approximately every 7–10 minutes depending on the speed at which the driver drove the car. The same simulation was used for all the subjects for uniformity. There was no constraint on the speed at which the subject chose to drive as the main focus of this work is on the driver's eyes and not on the performance of the driver.

The experiment was carried out with people who didn't wear glasses. This was done as the light from the camera was reflected by the glasses, as a result it could be very hard for the system to locate the eyes in this set-up.

In order to check that the system could give fatigue alerts the subjects closed their eyes for a few seconds from time to time. As the focus of the experiment was only on detection of the eyes, there were no other constraints for the test.

The images obtained from the video camera were used as input to the system, which consists of a Pentium Pro 200 MHz personal computer with a Matrox Genesis imaging board, which holds a Texas Instruments TMS320C80 DSP chip.

Chapter 6. DISCUSSION AND RESULTS

The images from all the 20 subjects were tested in the lab using the system as described above. Two pairs of templates for open and closed eyes were extracted by differencing and blob operations on the binarized difference image.

The system was tested on 20 different people of different skin color, with facial hair, and of different gender. We have monitored the system's response to different degrees of head orientation. The system rarely loses track of the eyes for small head movements. The system has a tolerance on head rotation of 45 degrees and on tilt up to 15 degrees. Furthermore, as the whole face region is used, the system doesn't need to relocate the search area unless there is a tracking error. The tracking error is also relatively lower as the whole face area is searched rather than just a small region localized around the eyes.

The system was able to detect blinks almost all the times and the system didn't give any false alarms in 19 out of the 20 cases considered.

The system was tested with eye templates of varying sizes for each subject. It was found that too small an eye template consisting of just the iris and the sclera for the open eyes and the same size of the template of the closed eyes didn't produce very good matches for some cases. There were mismatches especially in the case of closed eyes as the system finds any part of the skin region as the eye. Thus, there were misses due to incorrect matching with the facial hair for open eyes and other parts of the face for the closed eyes.

For too large a template, the misses were again high. So, an optimum template size consisting of the eyes with the eyelashes for the closed eyes and same size for the open eyes was used. The mismatch of the templates produced by the system with these templates is very minimal. The system can also detect the location of eyes correctly in most cases even with changes in lighting.

Another problem that was encountered was due to lighting. The room that was used was a semi-dark room. The room was enclosed on all sides by means of dark screens with the simulator screen in the front. However, there was light coming in from the nearby rooms and labs. As these were not direct sources, it was assumed that the only light was the light from the camera. However, there was also the light from the screen of the simulator reflected on the face of the driver. Further, the light from the projector placed at the top of the room also interfered with the camera light in some cases. In addition to these obstacles, a monitor kept at the back of the room, which simulated the graphics also, was a source of light.

The light reflected from the screen of the simulator had some effects on the performance of the system. The simulator consists of a stretch of road, floating in space and repeating again on itself, recursively. The simulator had to be reset every time the subject went off the road or at the end of the road stretch sometimes. Simulator reset resulted in temporary semi darkness due to switching off of the simulator, as the camera light was not very strong. As a result, the lighting was not uniform throughout. Due to this additional lighting from the simulator, the skin color of one subject was slightly different from the skin color model used for detecting the face. As a result, the face detection in case of this subject wasn't as accurate as in the case of the other subjects.

The system was generally tolerant to these noisy sources. The skin color of the individuals also affected the performance of the system slightly. This was due to the less ambient light available while recording the images. The only light sources were the light from the camera and the light reflected from the screen of the simulator. Due to the reduced lighting, the system is a bit sensitive to the skin color, and the color of iris matching with the color of hair.

The system performed very well for people with fair skin as their skin reflects more ambient light.

For one case when the color of the hair matched with the color of the iris, the system produced occasional mismatches finding the match of facial hair instead of the iris. A larger template including the eyebrows helped reduce the incorrect matches.

For cases when the subjects had darker skin, as the dark skin reflected less light and due to lower illumination, it was hard to distinguish the background from the face. Also in some of these cases, the color of the iris matched with the color of the hair. In such cases, it was very hard to find the correct matches. The problem also was accentuated by the reflection of the ceiling on their faces.

Due to the reduced lighting, the windshield of the car reflected the ceiling, which appeared as a white line on the face of the subjects in the image. This presented some problems in obtaining the templates for the eyes as well as matching. As a result, relatively low scores were obtained for these subjects.

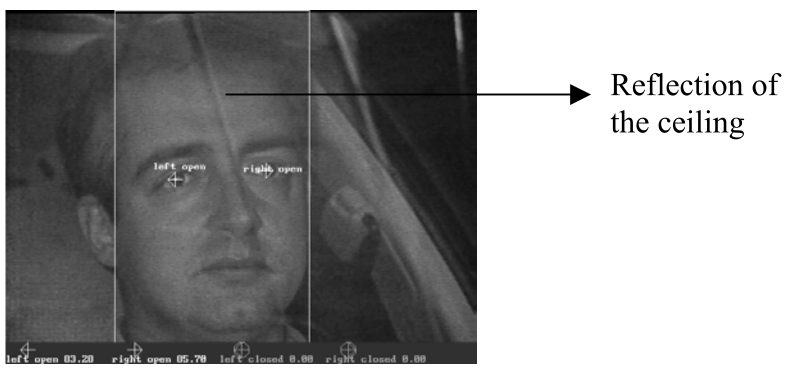

Figure 8. *Reflection of the ceiling on the subject's face.*

The reflection can be seen as a light strip of line cutting across the face. This was less apparent in this case as the skin color of the subject is fair. But the line looks much brighter and distinct as compared to the face in the case of subjects with darker skin as it hinders the detection of eyes in the image. The problem occurred due to the fact that the templates acquired previously were used for the ongoing detection of eyes and the templates were not updated dynamically. So when the face moved, the line seen may not appear at the same place on the face. If the templates acquired for example didn't consist of the line cutting over the eyes, when a person moved, the line might be over his eyes in a subsequent frame. In such a case, the match score for the eye would be lower. This is one of the main reasons for the low scores observed for a few cases in Table 1.

This problem was however less apparent in the case of subjects with fair skin compared to cases of the subjects with darker skin, as the line appeared brighter than the face. However, this problem is not of real concern as in the real-time application the camera is going to be mounted inside the car.

We tried to overcome the above problem by adding an additional light source pointing towards the face of the driver. This helped to fade out the reflection of the ceiling on the face and also brighten up the face of the driver against the background.

Table 1. *The performance of the system for the 20 subjects.*

Serial number	Average Open eye scores		Average Closed eye scores	
	Left eye open	Right eye open	Left eye closed	Right eye closed
1	70	70	70	70
2	70	70	70	70
3	75	75	78	78
4	75	75	85	85
5	60	60	70	70
6	75	75	85	85
7	87	65	90	85
8	85	83	83	83
9	75	75	73	73
10	73	73	80	80
11	67	67	85	85
12	68	50	85	68
13	85	85	85	85
14	85	85	85	85
15	90	90	90	90
16	80	80	80	80
17	70	70	70	70
18	77	77	77	77
19	20	20	20	20
20	50	55	58	65

The scores were averaged for the entire duration. As a result, the match scores in the case of "open eyes" are slightly lower than that of the "closed eye" match scores as these scores are also affected by the relative orientation of the face of the driver. The system failed in the case of subject "19" due to the reasons discussed in the previous pages.

Chapter 7. CONCLUSION AND FUTURE WORK

The system has been tested on 20 different subjects belonging to different races, gender, and having different skin color and facial hair. The system gave very good results for almost all of the cases. The system rarely loses track of eyes for relatively low head displacements. The system can also detect the blinks of the driver and gave no false alarms in 19 of the 20 cases studied. The failure in the one case was mainly due to the lighting conditions as discussed earlier.

The face of the driver was isolated using skin color information. The performance of the system could be improved if the foreground and the background information were isolated and used for face detection. Further, instead of using the entire face region for searching the eyes, a smaller search space consisting of a relatively smaller region around the eyes tracked in a previous stage could be used in the subsequent images. This was tried but the results produced by the system were not very satisfactory.

RERERENCES

[1.] Baluja, S. and Pomerleau, D. (1994). "Non-intrusive Gaze Tracking Using Artificial Neural Networks," Research Paper CMU-CS-94-102, School of Computer Science, Carnegie Mellon University: Pittsburgh, Penn., U.S.A

[2.] Brunelli, R. and Poggio, T. (1993). "Face Recognition: Features Versus Templates*," IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 15, No. 10, pp. 1042-1052.

[3.] Chow, G. and Li, X. (1993). "Towards a System for Automatic Feature Detection*," Pattern Recognition*, Vol. 26, No. 12, pp. 1739-1755.

[4.] Cox, I.J., Ghosn, J. and Yianilos, P.N. (1995). "Feature-Based Recognition Using Mixture-Distance," *NEC Research Institute, Technical Report 95-09.*

[5.] Craw, I., Ellis, H. and Lishman, J.R. (1987). "Automatic Extraction of Face-Features," *Pattern Recognition Letters, 5*, pp. 183-187.

[6.] De Silva, L.C., Aizawa, K. and Hatori, M. (1995). "Detection and Tracking of Facial Features by Using Edge Pixel Counting and Deformable Circular Template Matching*," IEICE Transaction on Information and Systems*, Vol. E78-D, No. 9, pp. 1195-1207, September.

[7.] Dijk, D.J. and Czeisler, C.A. (1994). "Paradoxal Timing of the Circadian Rhythm of Sleep Propensity Serves to Consolidate Sleep and Wakefulness in Humans," *Neuroscience Letters*, Vol. 166, pp. 63-68.

[8.] Dinges, D.F. "An Overview of Sleepiness and Accidents," *Journal of Sleep Research (supplement)*, *4*, pp. 1-11, 1995.

[9.] Huang, C.L. and Chen, C.W. (1992). "Human Facial Feature Extraction for Face Interpretation and Recognition," *Pattern Recognition*, Vol. 25, No. 12, pp. 1435-1444.

[10.] Hutchinson, R.A. (1990). "Development of an MLP Feature Location Technique Using Preprocessed Images," in *Proceedings of International Neural Networks Conference*, 1990.

[11.] Jochem, T.M., Pomerleau, D.A. and Thorpe, C.E. (1993). "MANIAC: A Next Generation Neurally Based Autonomous Road Follower," in *Proceedings of the International Conference on Intelligent Autonomous Systems* (IAS-3).

[12.] Kass, M., Witkin, A. and Terzopoulos, D. (1988). "Snakes: Active Contour Models," *International Journal of Computer Vision*, pp. 321-331.

[13.] Li, X. and Roeder, N. (1995). "Face Contour Extraction from Front-View Images," *Pattern Recognition*, Vol. 28, No. 8, pp. 1167-1179.

[14.] Nodine, C.F., Kundel, H.L., Toto, L.C. and Krupinksi, E.A. (1992). "Recording and Analyzing Eye-Position Data Using a Microcomputer Workstation," *Behavior Research Methods, Instruments and Computers, 24* (3), 475-584.

[15.] Pomerleau, D.A. (1991). "Efficient Training of Artificial Neural Networks for Autonomous Navigation," *Neural Computation* 3:1, Terrence Sejnowski (Ed).

[16.] Roeder, N. and Li, X. (1996). "Accuracy Analysis for Facial Feature Detection," *Pattern Recognition*, Vol. 29, No. 1, pp. 143-157.

[17.] Segawa, Y., Sakai, H., Endoh, T., Murakami, K., Toriu, T. and Koshimizu, H. (1996). "Face Recognition Through Hough Transform for Irises Extraction and Projection Procedures for Parts Localization," *Pacific Rim International Conference on Artificial Intelligence,* pp. 625-636.

[18.] Stringa, L. (1993). "Eyes Detection for Face Recognition," *Applied Artificial Intelligence*, No. 7, pp. 365-382.

[19.] Stiefelhagen, R., Yang, J. and Waibel, A. (1996). "A Model-Based Gaze-Tracking System," *International IEEE Joint Symposia on Intelligence and Systems*, pp. 304-310.

[20.] Tello, E.R. (1983). "Between Man and Machine," *BYTE*, September, pp. 288-293.

[21.] Tock, D. and Craw, I. (1995). "Tracking and Measuring Drivers' Eyes," *Real-Time Computer Vision*, pp. 71-89.

[22.] White, K.P., Hutchinson, T.E. and Carley, J.M. (1993). "Spatially Dynamic Calibration of an Eye-Tracking System," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 23, No. 4, pp. 1162-1168.

[23.] Wu, H., Chen, Q. and Yachida, M. (1996). "Facial Feature Extraction and Face Verification," *IEEE Proceedings of International Conference of Pattern Recognition,* pp. 484-488.

[24.] Xie, X., Sudhakar, R., and Zhuang, H. (1994). "On Improving Eye Features Extraction Using Deformable Templates," *Pattern Recognition*, Vol. 27, No. 6, pp. 791-799.

[25.] Xie, X., Sudhakar, R. and Zhuang, H. (1995). "Real-Time Eye Feature Tracking from a Video Image Sequence Using Kalman Filter," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 25, No. 12, pp. 1568-1577.

[26.] Yoo, T.W. and Oh, I.S. (1996). "Extraction of Face Region and Features Based on Chromatic Properties of Human Faces," *Pacific Rim International Conference on Artificial Intelligence,* pp. 637-645.