

Fusion Based Face Recognition Using Statistics of Shaded Subregions

ECE 533 Final Project

December 21, 2006

John A. Boehm

Introduction:

Face recognition is a growing field in image processing and machine learning with important and useful applications in surveillance, authorization and many other security applications. Recent advancements in the state of the art include using techniques operating on 3-D range data [1] and image capture under Near Infra-Red lighting [2] offer great promise for the future but can do little for helping to build a system today using existing databases of uncontrolled 2-D images. Uncontrolled images can come in the form of varying lighting, which introduces shadowing effects, pose variations that constrain a system to doing recognition on a mere fraction of the face, changes in expression that can include contortions of a face into an unrecognizable subject or poor quality, noisy images that simply do not contain enough information for recognition systems to operate. The most common of these problems is uncontrolled lighting.

In unprocessed images, changes between images of the same person under different lighting conditions are larger than those between two different people under the same lighting conditions [3]. Given the extensive number of images that exist, robust methods for pre-processing and recognition algorithms are clearly worthwhile topics for investigation. I have been working on the extensive set of images provided in the FRGC [4] database.

The Face Recognition Grand Challenge (FRGC) sponsored by NIST in 2005-06 was conceived in an attempt to improve the current state of the art of face recognition systems by an order of magnitude. Six different challenges were issued ranging from performing recognition on controlled, still images to using combinations of 2-D texture and 3-D range scan data. Of the six different challenges issued, experiment 4 is arguably the most challenging as it contains 2-D images of uncontrolled pose, expression and lighting variations in the data set. In an effort to standardize the metrics used for comparing algorithms and to insure reproducibility of results the FRGC offered the use of the Biometric Experiment Environment (BEE) as an operating platform.

The baseline recognition algorithm for experiment 4 provided with the BEE was only able to produce ~14% successful recognition (True Acceptance Rate or TAR) when the False Acceptance Rate (FAR) threshold was set at 0.1% (the standard used to judge all algorithms in the competition). My project deals with trying to improve the existing baseline algorithm by gathering statistics of sub-regions in each image, using the statistics to create weights for emphasizing the “more important regions” and generating a new matrix of similarity scores that improves the overall performance of the baseline algorithm.

Background:

The baseline algorithm first pre-processes the images in the database using meta-data provided for each image that provides the locations for critical landmarks, such as the corners of the eyes. This facilitates the rotation and translation of the image for proper

registration. It then extracts an oval shaped region containing just the face and performs histogram equalization on the region before storing it as a 1X19500 vector. A typical pre-processed controlled image is shown in **fig. 1** while **fig. 2** is a pre-processed uncontrolled image of the same person. The shading of the areas around the inside corners of the eyes as well as the regions below the nose and lips are quite evident.



The next module in the BEE is the biobox, which takes the normalized images from the pre-processing step as input and performs the actual recognition algorithm on them. The output of the biobox is an MxN matrix of similarity scores where M is the number of Query images (unknown identity) and N is the number of known Target images (known identity). Therefore, the i,j^{th} similarity score is the distance measure between the i^{th} Query image and the j^{th} Target image. The lower the score the closer the match is between the two. Finally, after normalizing the similarity matrix, a Receiver Operating Curve (ROC) is generated which plots the True Acceptance Rate (TAR) as a function of the False Acceptance Rate (FAR). The results of the baseline algorithm are shown in **fig. 3** and **fig. 4** below for both intra-semester and inter-semester recognition respectively. The standard used in the FRGC for comparing algorithms was the True Recognition Rate at a False Acceptance Rate threshold of 0.1%. Thus the score for the two curves below are 13.6% and 12.0% respectively.

