Face Recognition

Introduction

Humans often use faces to recognize individuals and advancements in computing capability over the past few decades now enable similar recognitions automatically. Early face recognition algorithms used simple geometric models, but the recognition process has now matured into a science of sophisticated mathematical representations and matching processes. Major advancements and initiatives in the past ten to fifteen years have propelled face recognition technology into the spotlight. Face recognition can be used for both verification and identification (open-set and closed-set).

History

Automated face recognition is a relatively new concept. Developed in the 1960s, the first semi-automated system for face recognition required the administrator to locate features (such as eyes, ears, nose, and mouth) on the photographs before it calculated distances and ratios to a common reference point, which were then compared to reference data. In the 1970s, Goldstein, Harmon, and Lesk[1] used 21 specific subjective markers such as hair color and lip thickness to automate the recognition. The problem with both of these early solutions was that the measurements and locations were manually computed. In 1988, Kirby and Sirovich applied principle component analysis, a standard linear algebra technique, to the face recognition problem. This was considered somewhat of a milestone as it showed that less than one hundred values were required to accurately code a suitably aligned and normalized face image.[2] In 1991, Turk and Pentland discovered that while using the eigenfaces techniques, the residual error could be used to detect faces in images[3] – a discovery that enabled reliable real-time automated face recognition systems. Although the approach was somewhat constrained by environmental factors, it nonetheless created significant interest in furthering development of automated face recognition technologies.[3] The technology first captured the public's attention from the media reaction to a trial implementation at the January 2001 Super Bowl, which captured surveillance images and compared them to a database of digital mugshots. This demonstration initiated much-needed analysis on how to use the technology to support national needs while being

considerate of the public's social and privacy concerns. Today, face recognition technology is being used to combat passport fraud, support law enforcement, identify missing children, and minimize benefit/identity fraud.

Predominant Approaches

There are two predominant approaches to the face recognition problem: geometric (feature based) and photometric (view based). As researcher interest in face recognition continued, many different algorithms were developed, three of which have been well studied in face recognition literature: Principal Components Analysis (PCA), Linear Discriminant Analysis (LDA), and Elastic Bunch Graph Matching (EBGM).

PCA: Principal Components Analysis (PCA)

PCA, commonly referred to as the use of eigenfaces, is the technique pioneered by Kirby and Sirivich in 1988. With PCA, the probe and gallery images must be the same size and must first be normalized to line up the eyes and mouth of the subjects within the images. The PCA approach is then used to reduce the dimension of the data by means of data compression basics[2] and reveals the most effective low dimensional structure of facial patterns. This reduction in dimensions removes information that is not useful[4] and precisely decomposes the face structure into orthogonal (uncorrelated) components known as eigenfaces. Each face image may be represented as a weighted sum (feature vector) of the eigenfaces, which are stored in a 1D array. A probe image is compared against a gallery image by measuring the distance between their respective feature vectors. The PCA approach typically requires the full frontal face to be presented each time; otherwise the image results in poor performance.[4] The primary advantage of this technique is that it can reduce the data needed to identify the individual to 1/1000[th] of the data presented.[5]
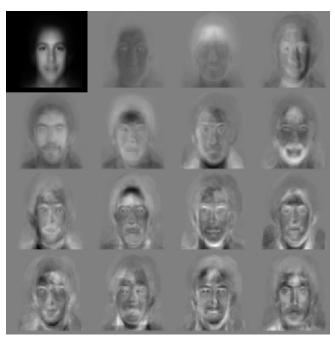
Figure 1: Standard Eigenfaces: Feature vectors are derived using eigenfaces.[6]

LDA: Linear Discriminant Analysis

LDA is a statistical approach for classifying samples of unknown classes based on training samples with known classes.[4] (Figure 2) This technique aims to maximize between-class (i.e., across users) variance and minimize within-class (i.e., within user) variance. In Figure 2 where each block represents a class, there are large variances between classes, but little variance within classes. When dealing with high dimensional face data, this technique faces the small sample size problem that arises where there are a small number of available training samples compared to the dimensionality of the sample space.[7]

Figure 2: Example of Six Classes Using LDA [8]

EBGM: Elastic Bunch Graph Matching

EBGM relies on the concept that real face images have many non-linear characteristics that are not addressed by the linear analysis methods discussed earlier, such as variations in illumination (outdoor lighting vs. indoor fluorescents), pose (standing straight vs. leaning over) and expression (smile vs. frown). A Gabor wavelet transform creates a dynamic link architecture that projects the face onto an elastic grid.[4] The Gabor jet is a node on the elastic grid, notated by circles on the image below, which describes the image behavior around a given pixel. It is the result of a convolution of the image with a Gabor filter, which is used to detect shapes and to extract features using image processing. [A convolution expresses the amount of overlap from functions, blending the functions together.] Recognition is based on the similarity of the Gabor filter response at each Gabor node.[4] This biologically-based method using Gabor filters is a process executed in the visual cortex of higher mammals. The difficulty with this method is the requirement of accurate landmark localization, which can sometimes be achieved by combining PCA and LDA methods.[4]
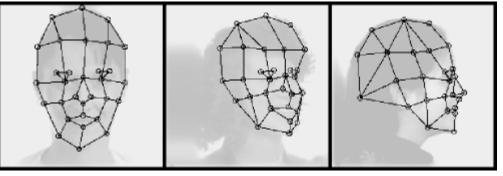

Figure 4: Elastic Bunch Map Graphing.[9]

United States Government Evaluations

The US Government has performed multiple evaluations to determine the capabilities and limitations of face recognition, and to encourage and direct future development. The FacE REcognition Technology (FERET) Evaluation, sponsored from 1993-1997 by the Defense Advanced Research Products Agency (DARPA),[10] was an effort to encourage the development of face recognition algorithms and technology by assessing the prototypes of face recognition systems. It propelled face recognition from its infancy to a market of commercial products.

The Face Recognition Vendor Tests (FRVT) were performed in 2000 and 2002, and another is planned for 2006. These evaluations built upon the work of FERET and coincided with the general onset of commercially available face recognition products. FRVT 2000[11] had two goals:

- Assess the capabilities of commercially available facial recognition systems; and

- Educate the biometrics community and the general public on how to properly present and analyze results.

FRVT 2002[12] was designed to measure technical progress since 2000, to evaluate performance on real-life large-scale databases, and to introduce new experiments to help better understand face recognition performance better. The FRVT 2002 included experiments with error bars, showing variances in performance as similar images were interchanged. Key FRVT 2002 results are:

- Given reasonable controlled indoor lighting, the current state of the art in face recognition is 90% verification at a 1% false accept rate.

- The use of morphable models, which maps a 2D image onto a 3D grid in an attempt to overcome lighting and pose variations, can significantly improve non-frontal face recognition.

- Watch list performance decreases as a function of gallery size – performance using smaller watch lists is better than performance using larger watch lists.

- In face recognition applications, accommodations should be made for demographic information since characteristics such as age and sex can significantly affect performance.

The goal of the Face Recognition Grand Challenge (FRGC) — the next step in the government development and evaluation process — is to promote and advance face recognition technology designed to support existing face recognition efforts of the US Government.[13] The FRGC will attempt to develop new face recognition techniques and develop prototype systems while increasing performance by an order of magnitude. The FRGC is open to face recognition researchers and developers in companies, academia, and research institutions. Soon after the completion of the FRGC, the Government will perform an in-depth assessment of face recognition — the FRVT 2006.

National Science and Technology Council (NSTC)

Committee on Technology

Committee on Homeland and National Security

Subcommittee on Biometrics

Standards Overview

Standardization is a vital portion of the advancement of the market and state of the art. Much work is being done at both the national and international standard organization levels to facilitate the interoperability and data interchange formats, which will help facilitate technology improvement on a standardized platform. The ANSI/INCITS (M1) 385-2004 and ISO SC37 19794-5 Face Recognition Data Interchange Format[14] are the major face recognition standards and address detailed human examination of face images, human verification of identity, and automated face identification and verification. These standards allow for interoperability among face recognition vendors.[15] The standards have established a defined frontal image[15] and are broken into subsections addressing full-frontal and token images. (A full-frontal image is defined as an image within five degrees from the center. A token image is defined by the location of the eyes.) These standards leave other images, such as semi-profile, undefined[15] but ensure that enrolled images will meet a quality standard needed for both automated face recognition and human inspection of face images.[14] Work is underway at both the national and international levels to update the standards for 3D face data. ANSI NIST ITL 1-2000 is also being updated to include more/better information for Type-10 face images. There is also related work at the international level to provide guidance to photographers on how to best capture face images for automated recognition. These standards also facilitate the use of face information in applications that have limited storage (e.g., passports, visas, driver's licenses).

Other standards, such as INCITS 398-2005 Common Biometric Exchange Formats Framework (CBEFF), deal specifically with the data elements used to describe the biometric data in a common way. The INCITS 358-2002 BioAPI Specification defines the Application Programming Interface and Service Provider Interface for a standard biometric technology interface. National and international standards organizations continue to work on the progression of standards in a direction that facilitates growth, advancement, and interoperability.

Summary

The computer-based face recognition industry has made much useful advancement in the past decade; however, the need for

higher accuracy systems remains. Through the determination and commitment of industry, government evaluations, and organized standards bodies, growth and progress will continue, raising the bar for face recognition technology.

Document References

[1] A. J. Goldstein, L. D. Harmon, and A. B. Lesk, "Identification of Human Faces," Proc. IEEE, May 1971, Vol. 59, No. 5, 748-760.

[2] L. Sirovich and M. Kirby, "A Low-Dimensional Procedure for the Characterization of Human Faces," J. Optical Soc. Am. A, 1987, Vol. 4, No.3, 519-524.

[3] M. A. Turk and A. P. Pentland, "Face Recognition Using Eigenfaces," Proc. IEEE, 1991, 586-591.

[4] D. Bolme, R. Beveridge, M. Teixeira, and B. Draper, "The CSU Face Identification Evaluation System: Its Purpose, Features and Structure," International Conference on Vision Systems, Graz, Austria, April 1-3, 2003. (Springer-Verlag) 304-311.

[5] "Eigenface Recognition" <http://et.wcu.edu/aidc/BioWebPages/eigenfaces.htm>.

[6] MIT Media Laboratory Vision and Modeling Group, "Photobook/Eigenfaces Demo" 25 July 2002 <http://vismod.media.mit.edu/vismod/demos/facerec/basic.html>.

[7] J. Lu, K.N. Plataniotis, and A.N. Venetsanopoulos, "Regularized Discriminant Analysis For the Small Sample Size Problem in Face Recognition," Pattern Recognition Letters, December 2003, Vol. 24, Issue 16: 3079-3087.

[8] Juwei Lu, "Boosting Linear Discriminant Analysis for Facial Recognition," 2002.

[9] Laurenz Wiskott, "Face Recognition by Elastic Bunch Graph Matching, " 24 April 1996 <http://www.neuroinformatik.ruhr-uni-bochum.de/ini/VDM/research/computerVision/graphMatching/identification/faceRecognition/contents.html>.

[10] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss, "The FERET Evaluation Methodology for Face-Recognition Algorithms," IEEE Transactions on PAMI, 2000, Vol. 22, No. 10: 1090-1104.

[11] D. M. Blackburn, J. M. Bone, and P. J. Phillips, "Facial Recognition Vendor Test 2000 Evaluation Report," February 2001 <http://www.frvt.org>.

[12] P. J. Phillips, P. Grother, R. J. Micheals, D. M. Blackburn, E. Tabassi, and J. M. Bone, "Face Recognition Vendor Test 2002 Overview and Summary," March 2003 <http://www.frvt.org>.

[13] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, "Overview of the Face Recognition Grand Challenge," Proc. Computer Vision and Pattern Recognition Conference, San Diego, 2005.

[14] "Information technology – Biometric data interchange formats – Part 5: Face image data." Documents ISO/IEC 19794-5:2005, 2004 <http://www.iso.org/>.

[15] "Information Technology - Face Recognition Format for Data Interchange," document 385-2004 ANSI INCITS, 2004 <http://www.incits.org/>.