© 2017 IJSRST | Volume 3 | Issue 3 | Print ISSN: 2395-6011 | Online ISSN: 2395-602X
Themed Section: Science and Technology

Pose Regularization Based Automatic Multi-View Face Recognition Method

Girishma Batra, Pramod Patil
Electronics and Telecommunication Department, Jhulelal Institute of Technology, Nagpur, Maharashtra, India

ABSTRACT

Face recognition techniques encounters difficulty in handling arbitrary poses variations. Various approaches have been worked upon for face recognition across pose variations, wherein many methods either require manual landmark annotations or assume the face poses to be known. These constraints prevent many face recognition systems from working automatically. The proposed work presents a new fully automatic multi-view face recognition method via 3D model based pose regularization, and extends existing face recognition systems into multi-view scenarios. Unlike previous pose normalization approaches, where non-frontal face images were transformed into frontal images, the proposed 3D model based pose regularization method generates synthetic target images to resemble the pose variations in query images. We should point out that generating non-frontal views from frontal face images is much easier and more accurate than recovering frontal views from non-frontal face images. This is because it is difficult to automatically detect accurate landmarks under large pose variations which are required to build a 3D face model. Additionally, since many areas of a face are significantly occluded under large pose variations, it is problematic to recover the frontal view for the occluded facial regions.The proposed work follows the view rendering based on 3D GEM but uses a simplified 3D Morphable Model [6]. Additionally, instead of aligning the synthetic target images and testing face images based on eye positions, we perform face alignment using Procrustes analysis under large pose variations. Moreover, our face matching method with blocked MLBP features provides better robustness against face illumination and expression variations. Finally, we show the expansibility of the proposed approach by replacing our MLBP based face matcher with two state- of-the-art face matching systems.

Keywords: Face recognition, pose regularization, feature point, 3D GEM, 3D Morphable Model

I. INTRODUCTION

The goal of automated face recognition (AFR) is to automatically recognize a person from digital images or video sequences containing his face. AFR has attracted substantial attention in the past decades due to its wide applications in real-world scenarios ranging from mobile phone authentication to surveillance. While AFR in controlled conditions, such as frontal or near-frontal poses, neutral expressions and near uniform illumination, has shown impressive performance, AFR in uncontrolled environments, such as arbitrary poses, non-uniform illumination, and partial occlusion, remains a challenging problem.

One typical application of AFR in uncontrolled conditions is identification or authorization of individuals with face images or videos captured by mobile devices, such as handheld terminals, mobile phones, or surveillance cameras. In these scenarios, there is a high possibility that the face images are captured without the cooperation of subjects. As a result, faces in the query images can be of arbitrary poses. Despite the potential value of non-frontal face images in forensic applications, the arbitrary pose variations have become one of the primary stumbling blocks for most existing systems to perform face recognition automatically.

The increasing ubiquity of surveillance imaging devices offers a promising avenue to combat acts of terrorism and crime. Tragic events such as the Boston bombings in 2013 and the 2011 riots in London have drawn attention to the potential role that surveillance cameras and video-based face recognition can play in identifying perpetrators of such criminal acts. Unfortunately, identification technology has not quite lived up to expectations in such instances. A lack of robustness to classic challenges in pose, illumination, and expression are the prevailing explanations as to why face recognition has limited performance in these unconstrained scenarios. However, equally important is the inability of state-of-the-art technology to leverage the additional temporal information available in sequences of face images (i.e., videos). In order for researchers to improve upon these limitations, we first need to determine the accuracy of deploy- able technology (i.e., a COTS face matcher) in matching unconstrained faces in video data. Without such a baseline, we cannot determine if meaningful progress is being made in video-based face recognition.

A person's face contains important clues for social interaction, providing a wide variety of useful information, including the person's identity, age, gender, race, expression, etc. Over the past 50 years, significant advances have been made on extracting discriminative features in a face image to determine the subject's identity [1]. In recent years, several applications have emerged that make use of demographic or soft biometric traits (e.g., age, gender, and race). These applications include access control, re-identification in surveillance videos, law enforcement, integrity of face images in social media, intelligent advertising, and human-computer interaction. As a result, studies on the exploration of various attributes (other than the identity) embodied in a face image, such as age, gender, and race, have drawn increasing attention [2]–[5]. While state- of-the-art demographic estimation methods are able to attain a mean absolute error (MAE) of about 4 years for age estimation, and more than 95% accuracy for gender and race classifications, most of these studies have utilized face images captured in rather controlled sensing and cooperative subject scenarios [6], such as the MORPH [7] face database. However, in many of the applications mentioned above, especially video surveillance, the available face images of a person of interest are most likely to be captured under unconstrained and un- cooperative scenarios (see Fig. 1). To close the gap between the applicability of published methods and the requirements of real world applications, researchers have attempted to develop new approaches for demographic estimation that are robust to unconstrained face images [8]–[13].

As face recognition applications progress from constrained imaging and cooperative subjects (e.g., identity card reduplication) to unconstrained imaging scenarios with uncooperative subjects (e.g., watch list monitoring), a lack of guidance exists with respect to optimal approaches for integrating face recognition algorithms into large-scale applications of interest. In this work we explore the problem of identifying a person of interest given a variety of information sources about the person (face image, surveillance video, face sketch, 3D face model, and demographic information) in both closed set and open set identification modes.

Identifying a person based on unconstrained face images is an increasingly prevalent task for law enforcement and intelligence agencies. In general, these applications seek to determine the identity of a subject based on one or more probe images or videos, where a top 200 ranked list retrieved from the gallery (for example) may suffice for analysts (or forensic examiners) to identify the subject [1]. In many cases, such a forensic identification is performed when multiple face images and/or a face track (i.e., a sequence of cropped face images which can be assumed to be of the same person) from a video of a person of interest are available. For example, in investigative scenarios, multiple face images of an unknown subject often arise from an initial clustering of visual evidence, such as a network of surveillance cameras, the contents of a seized hard drive, or from open source intelligence (e.g., social networks). In turn, these probe images are searched against large-scale face repositories, such as mug shot or identity card databases.

While other routine, but high value, crimes such as armed robberies, kidnappings, and acts of violence require similar identifications, only a fraction of the manual resources are available to solve these crimes. Thus, it is paramount for face recognition researchers and practitioners to have a firm understanding of optimal strategies for combining multiple sources of face information, collectively called face media, available to identify the person of interest.

While forensic identification is focused on human-driven queries, several emerging applications of face recognition technology exist where it is neither practical nor economical for a human to have a high degree of intervention with the automatic face recognition system. One such example is watch list identification from surveillance cameras, where a list of persons of interest is continuously searched against streaming videos. Termed as open set recognition, these challenging applications will likely have better success as unconstrained face recognition algorithms continue to develop and mature.

Face recognition performances have shown computer vision capabilities to surpass those of humans. Rather than signaling the end of face recognition research, these results have led to a redefinition of the problem, shifting attention from highly regulated, controlled image settings to faces captured in unconstrained conditions (a.k.a., “in the wild”).

This change of focus, from constrained to unconstrained images, has toppled recognition rates (see, e.g., the original results on the Labeled Faces in the Wild benchmark, published in the same year). This drop was not surprising: Unconstrained photos of faces represented a myriad of new challenges, including changing expressions, occlusions, varying lighting, and non-frontal, often extreme poses. Yet in recent years recognition performance has gradually improved to the point where once again claims are being made for super-human face recognition capabilities.

During the past decade, face recognition has attracted much attention due to its great potential value in real world applications, such as access control, identity verification and video surveillance. However, in unconstrained environment the performance of face recognition always drops significantly because of large variations caused by pose, illumination, expression, and occlusion and so on. Among them pose and expression have always been important challenges because they can dramatically increase intra-person variances, sometimes even exceeding inter-person variances. To deal with the two challenges, many promising works have been developed, which can be divided into two categories: feature level normalization and image level normalization.

Face recognition, frequently performed unconsciously by humans has achieved a great deal of attention from the academic and industrial communities during the past two decades [1]. Face recognition aims at identifying or verifying a person’s identity by matching an input face image against the known faces in a database. Two basic face recognition categories are: (a) face identification and (b) face verification [2]. In face identification, a probe (test) image of an unknown individual is identified by comparing the image with an image gallery (training) of the known individuals. The identification scenario is also known as one-to-many (1:N) matching. In face verification a query face image is compared with only the image of a claimed identity. Alternatively, verification is the process of determining a person’s claimed identity. Face verification scenario is also known as one-to-one (1:1) matching. Fig. 1 shows the general procedure of a face recognition system. Initially, facial features are calculated and stored for the gallery (training) images. Later, these features are compared with the features of the probe (test) image and a similarity metric called score is computed for a given comparison.

In many real-world applications, human faces are captured in unconstrained environments. The performance of the face recognition algorithms drops rapidly due to various factors such as facial expressions, background clutter, surgery, and hairstyle [3]. Moreover, occlusion, non-uniform illuminations (shadows, underexposure, and overexposure), pose, and ageing significantly degrade the face recognition accuracy [4]. Owing to the aforementioned factors, face recognition has remained a challenging problem in pattern recognition. Owing to the rapid increase in assassinations and violence in recent times, face recognition systems demand even more attention in terms of accuracy and robustness when used in various domains such as forensic applications, access control, and security in public places. In such applications, the robustness of the system plays an important role [5]. Many databases of facial images [two-dimensional (2D)/3D], for example, [6–9] have been developed to test the accuracy of face recognition algorithms. Each database is designed to test a specific facial aspect such as pose, illumination, low resolution (LR), expression, and occlusion. Face-pose and image resolution are the two important factors that seriously challenge most of the developed face recognition algorithms.

II. RELATED WORK

Many researchers have carried out work in this domain.In this section we bring out the gist of the work carried out. Koichiro Niinuma, Hu Han, and Anil K. Jain in their work " Automatic Multi-view Face Recognition via 3D Model Based Pose Regularization" have proposed a fully automatic method for multiview face recognition. A 3D model is built from each frontal target face image, which is used to generate synthetic target face images. The pose of a query face image is also estimated using a multi-view face detector so that the synthetic target face images can be generated to resemble the pose variation of a query face image. Procrustes analysis is then applied to align the synthetic target images and the query image, and block based MLBP features are extracted for face matching. Experimental results on two public-domain databases (Color FERET and PubFig), and a Mobile face database collected using mobile phones show that the proposed approach outperforms two state-of-theart face matchers (FaceVACS and MKD-SRC) in automatic multi-view face recognition. The proposed approach can also be easily extended to leverage existing face recognition systems for automatic multi-view face recognition.
Lacey Best-Rowden , Brendan Klare , Joshua Klontz , and Anil K. Jain in their work "Video-to-Video Face Matching: Establishing a Baseline for Unconstrained Face Recognition" published in IEEE 2013 have studied the problem of video-to-video face matching in order to gain a better understanding of state-of-the-art recognition accuracies of COTS face matchers. Through studying the video-based face recognition problem, the methods proposed in this paper can be readily applied in any operational setting using existing COTS matchers (as opposed to dedicated video-based algorithms). Thus, the results provided in this paper offer researchers and practitioners a better understanding of how accurately video face data can be recognized using off-the-shelf technology.

The contributions of this paper can be summarized as follows. (i) A framework for applying COTS face recognition algorithms to video-based data is provided. (ii) The impact and consistency of different match score fusion rules are studied for several commercial matchers, providing guidance on how to best consolidate frame by frame face match scores across the video. (iii) The performance of COTS algorithms is studied with respect to quality-based key-frame subset selection. (iv) An order of magnitude decrease in error rates is achieved on the YouTube Faces Database, respective to the best accuracy currently re- ported in the literature.

Hu Han and Anil K. Jain in their work "Age, Gender and Race Estimation from Unconstrained Face Images" have presented an integrated framework for age, gender, and race estimation from unconstrained face images. The approach involves (i) face normalization consisting of pose and photometric corrections of an input face image, (ii) feature extraction from the normalized face image, including both the central face region and the surrounding contextual information (e.g., facial shape, ears, and hair style), and (iii) age group (or exact age), gender, and race estimation using Support Vector Machines (SVM).

Unlike previous studies where low-resolution face images (e.g., interpupillary distance (IPD) smaller than 24 pixels) were excluded from their evaluations [11], [13], we evaluate the proposed approach using the entire Images of Groups database [8] containing 28,231 face images. Additionally, we use an extended version (LFW+) of the public-domain LFW database containing 15,699 face images to perform age, gender, and race estimation.

The main contribution of this paper is to estimate age, gender and race from unconstrained face images. Most of the previous work focused only on MORPH and FG-NET databases. This is the first paper that gives complete and de- tailed results on age (EXACT age and age group), gender and race estimation for the LFW database. Novelty of this work consists of (i) cascade face normalization, and (ii) utilizing face context information. Face normalization handles pose and illumination variations in unconstrained faces. Additional contributions: (i) LFW+ database with EXACT age, gender, and race for each subject, and (ii) demographic estimation under cross-database scenarios.

Lacey Best-Rowden, Hu Han, Charles Otto, Brendan Klare, and Anil K. Jain in their work "Unconstrained Face Recognition: Identifying a Person of Interest from a Media Collection" examined the use of commercial off the shelf (COTS) face recognition systems with respect to the aforementioned challenges in large-scale unconstrained face recognition scenarios. First, the efficacy of forensic identification is explored by combining two public-domain unconstrained face databases, Labeled Faces in the Wild (LFW) [2] and YouTube Faces (YTF) [3], to create sets of multiple probe images and videos to be matched against a gallery consisting of a single image for each subject. To replicate forensic identification scenarios, we further populate our gallery with one million operational mug shot images from the Pinellas County Sheriff's Office (PCSO) database. Using this data, researchers were able to examine how to boost the likelihood of face identification through different fusion schemes, incorporation of 3D face models and hand drawn sketches, and methods for selecting the highest quality video frames. Researchers interested in improving forensic identification accuracy can use this competitive baseline (on public-domain databases LFW and YTF) to provide more objectivity towards such goals.

Most of the work on unconstrained face recognition using the LFW and YTF databases has been reported in verification scenarios [6], [7]. However, in forensic investigations, it is the identification mode that is of interest, especially the open-set identification scenario where the person of interest may not be present in legacy face databases.

The contributions of this work are summarized as follows:

- ✓ It has been shown for the first time, how a collection of face media (image(s), video(s), 3D model(s), demographic data, and sketch) can be used to mitigate the challenges associated with unconstrained face recognition (uncooperative subjects, unconstrained imaging conditions) and boost recognition accuracy.
- ✓ Unlike previous studies that report results in verification mode, results have been presented for both open set and closed set identifications which are the norm in identifying persons of interest in forensic and watch list scenarios.
- ✓ It presents effective face quality measures to determine when the fusion of information sources will help boost identification accuracy. The quality measures are also used to assign weights to different media sources in fusion schemes.
- ✓ To demonstrate the effectiveness of media-as-input for the difficult problem of unconstrained face recognition, researchers have utilized a state of the art COTS face matcher and a separate COTS 3D face modeler, namely the Aureus 3D SDK provided by Cyber Extruder. Face sketches were drawn by forensic sketch artists who generated the sketch after viewing low quality videos. In the absence of demographic data for LFW and YTF databases, crowd sourcing has been used to obtain the estimates of gender and race. The above strategy allows to show the contribution of various media components as we incrementally add them as input to the face matching system.
- ✓ Pose-corrected versions of all face images in the LFW database, pose-corrected video frames from the YTF database, forensic sketches, and experimental protocols used in this paper have been made publicly available.

Tal Hassner, Shai Harel, Eran Paz, Roee Enbar in their work " Effective Face Frontalization in Unconstrained Images"

Have attempted to estimate a rough approximation for the 3D surface of the face and use this surface to generate the new views. Although appealing, this approach relies on accurate localization of facial feature points and does not guarantee that the same alignment (frontalization) will be applied to different images of the same face. Thus, different images of the same person may well be aligned differently, preventing their features from being accurately compared.

They have proposed a simple alternative approach of using a single, unmodified 3D reference for all query faces in order to produce frontalized views. Ignoring individual differences in facial shapes may be counter-intuitive – indeed, previous work has emphasized its importance – however, qualitative examples throughout this paper show that any impact this has on facial appearances is typically negligible. In fact, faces remain easily recognizable despite this approximation. More importantly, the frontalized faces are aggressively aligned thereby improving performances over previous alignment methods. These claims are verified by showing elevated face verification results on the LFW benchmark and gender classification accuracy on the Adience benchmark, obtained using our frontalized faces.

Xiangyu Zhu Zhen Lei, Junjie Yan Dong Yi Stan Z. Li In their work " High-Fidelity Pose and Expression Normalization for Face Recognition in the Wild" have presented a pose and expression normalization method to recover the canonical-view, expression- free image with "high fidelity", which indicates preserving the face appearance with little artifact and information loss. The contributions are as follows: Firstly, a "land- mark marching" assumption is made to describe the movement of 3D landmarks across poses and propose a landmark based pose adaptive 3DMM fitting method. Secondly, an identity preserving normalization is proposed by meshing the whole image into a 3D object and normalizing it with 3D transformations. Finally, we propose a "Trend Fitting and Detail Filling" method to fill the in- visible region with Poisson editing, leading to smooth and natural normalization result. Based on the well developed landmark detector, the entire normalization system does not contain any learning procedure, leading to good generalization performance to different environments.

Zahid Mahmood Tauseef Ali, Samee U. Khan in their work "Effects of pose and image resolution on automatic face recognition" have reviewed the recent advances in face recognition and presented comparative study of three baseline face recognition algorithms. Face recognition algorithms studied in this work are: PCA, AdaBoost with LDA as a weak learner, and the LBP. The main goal of the study was to explore the robustness of each of these face recognition algorithms with respect to variation in pose and image resolution. Images from multi-PIE database were used for evaluation. For experimental setup, one frontal mug-shot was used in gallery while four different pose images are used as probe. For face size of 231 × 251 down to 30 × 30 pixels, the PCA-based algorithm was found to be more accurate in classifying the four poses followed by AdaBoost and the LBP algorithm. For LR images of size 20 × 20, 10 × 10, and 5 × 5 pixels, the AdaBoost-based algorithm surpassed the PCA and the LBP. The LBP-based face recognition algorithm failed to recognize face images of size 20 × 20 pixels and lower. For frontal face, AdaBoost-based algorithm exhibited 100% classification accuracy for size of 10 × 10 and 5 × 5 pixels. A major finding of the research was that LR images of the multi-PIE database do not affect the classification accuracy of AdaBoost-based face recognition algorithm.

Jae Young Choi, Yong Man Ro and Konstantinos N. Plataniotis in their work " Color Local Texture Features for Color Face Recognition" have proposed the first so-called color local texture features. Specifically, two effective color local texture features have been developed , i.e., color local Gabor wavelets (CLGWs) [15] and color LBP (CLBP) [16], both of which are able to encode the discriminative features derived from spatiochromatic texture patterns of different spectral channels (or bands) within a certain local region. In addition, to make full use of both color and texture information, the opponent color texture features that capture the texture patterns of spatial interactions between spectral bands are incorporated into the generation of CLGW and CLBP. This allows for ac-quiring more discriminative color local texture features, as compared with conventional grayscale texture features, for improving FR performance.

Comparative and extensive experiments have been conducted to investigate the effectiveness of proposed color local texture features. For this, five public face databases (DB), i.e., CMU-PIE [17], Color FERET [18], XM2VTSDB [19], SCface [20], and FRGC 2.0 [21, 22], are used. The experimental results show that the FR approach using our color local texture features achieves even better FR performance than the FR approach relying only on color or texture information.
designations.

III. PROPOSED WORK

The proposed method presents a new fully automatic multi-view face recognition method via 3D model based pose regularization, and extends existing face recognition systems into multi-view scenarios. Fig. 1 illustrates the proposed approach which consists of two main modules:

(i) Pose regularization based on 3D model, and
(ii) Face matching with block based multi-scale LBP (MLBP) features.

Unlike previous pose normalization approaches, where non-frontal face images were transformed into frontal images, the proposed 3D model based pose regularization method generates synthetic target images to resemble the pose variations in query images. We should point out that generating non-frontal views from frontal face images is much easier and more accurate than recovering frontal views from non-frontal face images. This is because it is difficult to automatically detect accurate landmarks under large pose variations which are required to build a 3D face model. Additionally, since many areas of a face are significantly occluded under large pose variations, it is problematic to recover the frontal view for the occluded facial regions.
The proposed pose regularization approach is similar to the novel view rendering based on 3D GEM, but the proposed method uses a simplified 3D Morphable Model [6]. Additionally, instead of aligning the synthetic target images and testing face images based on eye positions, we perform face alignment using Procrustes analysis under large pose variations. Moreover, our face matching method with blocked MLBP features provides better robustness against face illumination and expression variations. Finally, we show the expansibility of the proposed approach by replacing our MLBP based face matcher with two state- of-the-art face matching systems.

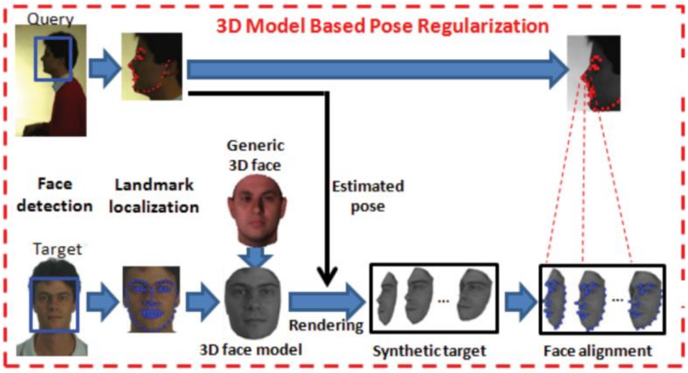

Figure 1. The Proposed Approach

A) 3D face modeling

In this work, we will utilize a simplified 3D Morphable Model [6] without the texture fitting due to its robustness and computational efficiency. We derive our 3D shape model from the USF Human ID 3-D database [2], which includes 3D face shape and texture of 100 subjects captured with a 3D scanner. The original 3D face includes 75,972 vertices, but for efficient computation, we interactively select 76 vertices based on the 76 keypoints defined in an open source Active Shape Model (Stasm). Given the 100 3D faces, the 3D shape of a new face can be represented using a PCA model.

A 2D face image is a projection of a 3D face onto a 2D plane under a set of transformations such as translation, rotation, scaling, and projection. Based on such a face imaging process, the shape of a 3D face can be recovered from its 2D projection (facial landmarks in a 2D face image) by minimizing the cost function.

B) Face alignment

By building a 3D model and generating synthetic target face images to resemble the poses of query images, we are able to reduce the pose disparity between them. However, face alignment is still necessary for the following feature extraction and face matching steps. Holistic face alignment based on two eyes (e.g. Inter-Pupil Distance (IPD)) has been a widely used approach for frontal or near-frontal face images. However, IPD based face alignment be- comes problematic for non-frontal poses. Under large pose variations one of the two eyes is often not visible, and even when both eyes are visible in non-frontal images, IPD based alignment can lead to an artificial increase in the overall size of the face image.

In our approach, we apply Procrustes analysis [11] to align the synthetic target images and a query image based on the facial landmarks from a 3D face model and keypoints that are detected by MTSPM. Although the numbers of key- points defined in a 3D face model and MTSPM are different, the keypoint sequence in each model is fixed. This makes it possible for us to manually establish the keypoint correspondence between two models. We have manually identified 19 landmarks in MTSPM that have corresponding landmarks in a 3D model. The Procrustes analysis is performed based on 19 corresponding landmark pairs.

C) Face matching

Given the aligned synthetic target face images and a query face image, we extract MLBP features for face matching. In our experiments, we use MLBP features which are a concatenation of LBP histograms with 8 neighbors sampled at different radii $R = \{1; 3; 5; 7\}$. We first divide a holistic face image (256×192) into 768 sub-regions (8×8 non-overlapped blocks). Then, MLBP features are extracted from individual blocks and concatenated together to represent a face.

Given two MLBP histograms x and y with n dimensions which are extracted from two face images, chi-squared distance $\chi2$ is calculated as a measure of similarity between two face images. The final distance between a target and a query is calculated by finding the minimum of these distances..

IV.RESULT AND DISCUSSION

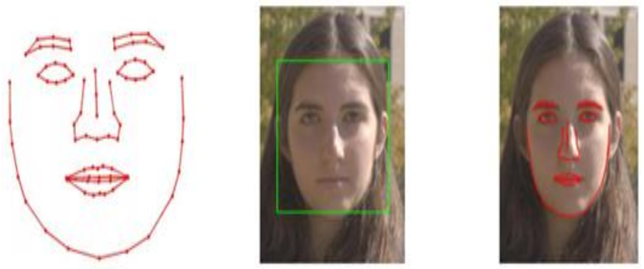

Figure 2. Face detection and feature point

Figure 2 shows the results of the face detection process based on the Viola-Jones framework. After the face has been detected (indicated by the green box), the active shape model (ASM) feature points are located using the Cootes method. These points are shown in red.

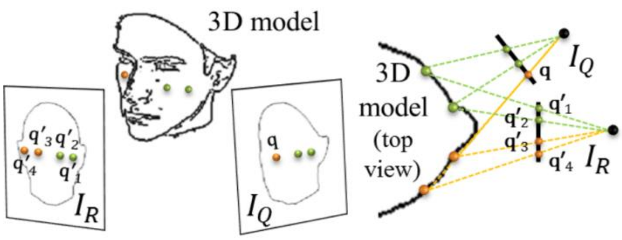

Figure 3. Face frontalization

Figure 3 depicts the process of face frontalization using the ASM feature points, which are superimposed on the morphable 3D face generic model proposed by Zhu-Ramanan. The frontalization is a result of altering the projection of the face feature points to a plane that is perpendicular to the 3D face model.

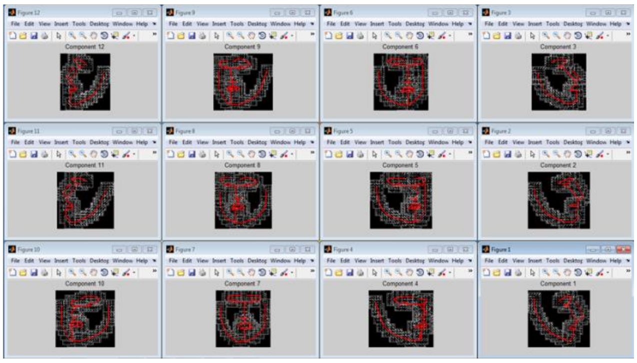

Figure 4. Face rotation Templates

Figure 4 shows the components that make up the morphable 3D face model proposed by Zhu Ramanan. The ASM face features in red are shown superimposed on the morphable 3D model.

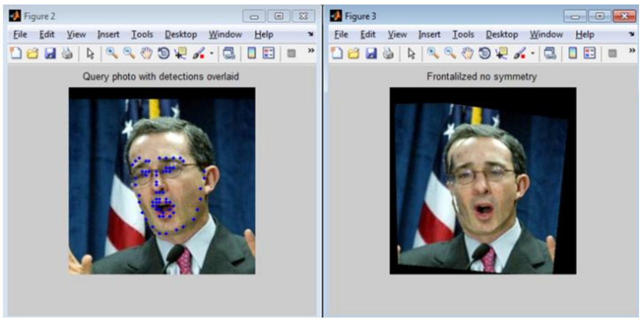

Figure 5. feature_points_and_frontalized_face

Figure 4 shows the result of the frontalization process on a sample image taken from the labeled faces in the wild (LFW) dataset. The blue points denote the ASM feature points proposed by Cootes. The output image is a result of rotation of the 3D morphable model into the perpendicular plane.

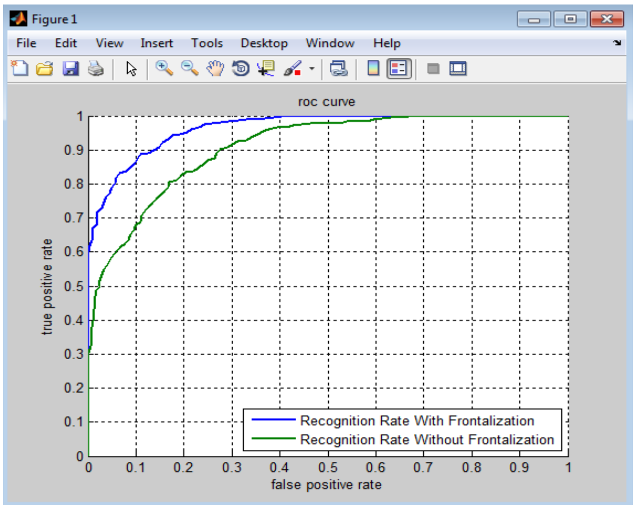

Figure 6. receiver_operator_curve

Figure 5 is the receiver operator curve for the detection rates achieved using frontalization prior to face recognition (the curve in blue). This is compared with the recognition results achieved without frontalization (the curve in green). Recognition with frontalization achieves a higher true positive rate for a lower false rate compared to recognition without frontalization. The measured difference is between one and two basis points.

V. CONCLUSION

In this work we have been able to present a system which carries out face identification of persons of interest in unconstrained imaging scenarios with uncooperative subjects. Given a face media collection of a person of interest (i.e., face images or video clips, 3D face models built from image(s) or video frame(s), face sketch, and demographic information), we have been able to demonstrate an incremental improvement in the identification accuracy of a face matching system. It is expected that this work may be a good supplement to the field of forensic investigations and watch list operations.

VI. REFERENCES

[1]. Zahid Mahmood, Tauseef Ali, Samee U. Khan, Effects of pose and image resolution on automatic face recognition, IET Journal of Biometrics, 2016.

[2]. Tal Hassner, Shai Harel, Eran Paz, Roee Enbar, Effective Face Frontalization in Unconstrained Images, IEEE 2015.

[3]. Xiangyu Zhu Zhen Lei Junjie Yan Dong Yi Stan Z. Li, Pose and Expression Normalization for Face Recognition in the Wild, IEEE Transactions on Biometrics 2015.

[4]. Hu Han, Member, IEEE and Anil K. Jain, Fellow, IEEE, Age, Gender and Race Estimation from Unconstrained Face Images, IEEE 2014.

[5]. Lacey Best-Rowden, Hu Han, Member, IEEE, Charles Otto, Brendan Klare, Member, IEEE, and Anil K. Jain, Fellow, IEEE, Unconstrained Face Recognition: Identifying a Person of Interest from a Media Collection, IEEE Transactions on Information Forensics and Security, 2014.

[6]. Koichiro Niinuma, Hu Han, and Anil K. Jain, Automatic Multi-view Face Recognition via 3D Model Based Pose Regularization, IEEE 6th International Conference on Biometrics USA, 2013.

[7]. Lacey Best-Rowden , Brendan Klare , Joshua Klontz , Anil K. Jain, Establishing a Baseline for Unconstrained Face Recognition, IEEE Transactions on Biometrics 2013.

[8]. S. R. Arashloo and J. Kittler. Energy normalization for pose-invariant face recognition based on MRF model image matching. IEEE Trans. PAMI, 33(6):1274–1280, 2011.

[9]. A. Asthana, T. K. Marks, M. J. Jones, K. H. Tieu, and M. V. Rohith. Fully automatic pose invariant face recognition via 3D pose normalization. Proc. ICCV, pages 937–944, 2011.

[10]. C. D. Castillo and D. W. Jacobs. Wide-baseline stereo for face recognition with large pose variation. Proc. IEEE CVPR, pages 537–544, 2011.

[11]. Liu, Z., Lv, S., Lv, Y., et al.: ‘A natural visible and infrared facial expression database for expression recognition and emotion inference’, IEEE Trans. Multimedia, 2010, 12, (7), pp. 682–691

[12]. A. B. Ashraf, S. Lucey, and T. Chen. Learning patch correspondences for improved viewpoint invariant face recognition. Proc. IEEE CVPR, pages 1–8, 2008.

[13]. Huang, G., Ramesh, M., Berg, T., et al.: ‘Labeled faces in the wild: a database for studying face recognition in unconstrained environments’. Technical Report, 07–49, University of Massachusetts, Amherst, 2007.

[14]. Xiujuan, C., Shiguang, S., Xilin, C., et al.: ‘Locally linear regression for pose-invariant face recognition’, IEEE Trans. Image Process., 2007, 16, (7),pp. 1716–1725

[15]. Phillips, P.J., Flynn, P.J., Scruggs, T., et al.: ‘Overview of the face recognition grand challenge’, Comput. Vis. Pattern Recognit., 2005, (1), pp. 947–954

[16]. Viola, P., Jones, M.: ‘Robust real-time face detection’, Int. J. Comput. Vis, 2004, 57, pp. 137–154

[17]. V. Blanz and T. Vetter. Face recognition based on fitting a 3D morphable model. IEEE Trans. PAMI, 25(9):1063–1074, 2003.

[18]. Volker, B., Thomas, V.: ‘Face recognition based on fitting a 3-D morphable model’, IEEE Trans. Pattern Anal. Mach. Intell., 2003, 25, (9), pp. 1063–1074

[19]. Phillips, J., Wechsler, H., Huang, J.S., et al.: ‘The FERET database and evaluation procedure for face recognition algorithms’, Image Vis. Comput., 1998, pp. 295–306

[20]. Turk, M., Pentland, A.: ‘Eigenfaces for recognition’, J. Cogn. Neurosci., 1991, 3, (1), pp. 71–86

[21]. FaceVACS Software Developer Kit Cognitec Systems GmbH, http://www.cognitec-systems.de/.

[22]. USF DARPA HumanID 3D Face Database. Sudeep Sarkar, University of South Florida.