

POSE AND ILLUMINATION COMPENSATION FOR 3D FACE RECOGNITION

Sotiris Malassiotis and Michael G. Strintzis

Informatics and Telematics Institute, Thessaloniki, Greece
e-mail: malasiot@iti.gr

ABSTRACT

The paper describes a face recognition system using a combination of color and depth images. To cope with illumination and pose variations 3D information is used for the normalization of the input images. The proposed pose compensation algorithm is based on a robust 3D face detection and pose estimation technique, while illumination compensation exploits depth data to recover the illumination of the scene and relight the image under frontal lighting. When normalized images, depicting upright orientation and frontal lighting, are used for classification significantly high recognition rates are achieved, as demonstrated on a face database with more than 2000 images.

1. INTRODUCTION

Recent public face recognition tests demonstrated that the accuracy of state-of-the-art algorithms degrades significantly for images exhibiting pose and illumination variations. Current research efforts strive to achieve insensitivity to such variations.

The paper describes and evaluates a complete face identification system using a combination of 2D color and 3D range images captured in real-time. We present several novel techniques which are capable, taking as input a pair of 2D and 3D images, to produce a pair of normalized images depicting frontal pose and illumination. The efficiency and robustness of the proposed system is demonstrated on a data set of significant size and compared with state-of-the-art compensation techniques.

Although the 3D structure of the human face conveys important discriminatory information only a few techniques have been proposed employing range images. This is mainly due to the high cost of available 3D digitizers and the fact that they do not operate in real time (e.g. time of flight laser scanners) or produce inaccurate depth information (e.g. stereo vision). The work presented in this paper is partly motivated by the recent development of novel low cost 3D of low-cost

sensors that are capable of real-time 3D acquisition [1]. A common approach adopted towards 3D face recognition is based on the extraction of 3D facial features by means of differential geometry techniques [2–4]. A few techniques [5,6] also employ grayscale images but mainly for augmenting the detection of features such as the eyes that are harder to detect on the range image. Although feature-based techniques are robust to pose variations they rely on accurate 3D maps of faces, usually extracted by expensive off-line 3D scanners. Thus their applicability to real-world applications with highly noisy data is questionable. The recognition rates claimed by the above techniques were estimated using databases of limited size and without significant variations of the faces. Only recently [7] conducted an experiment with a database of significant size (275 persons) containing both grayscale and range images, and produced comparative results of face identification using eigenfaces for 2D, 3D and their combination and for varying image quality. This test however considered only frontal images captured under constant illumination conditions. For this work we have recorded a face database containing several appearance variations. These variations are compensated before reaching the classifier, thus leading to high recognition rates.

2. ACQUISITION OF 3D DATA

The proposed system is based on real-time quasi-synchronous color and 3D image acquisition based on the color structured-light approach [1]. The sensor is based on low cost devices, an off-the-shelf CCTV-color camera and a standard slide projector. The average depth accuracy of the system optimized for an access control application is about $0.5mm$. The spatial resolution of the range images is approximately equal to the color camera resolution.

Using the above setup a face database was recorded. For each subject several images depicting different appearance variations were acquired: three facial expressions, three types of illumination (left/right side spot lights and overhead light), two pose variations (± 20 degrees), two images with and without glasses, and three frontal images. The database contains 20 persons and 2 recordings with time lapse between

This work is funded by research project BioSec IST-2002-001766 (Biometrics and Security, <http://www.biosec.org>), under Information Society Technologies (IST) priority of the 6th Framework Programme of the European Community.

each recording session about 10 days (2200 image pairs).

3. POSE COMPENSATION

The aim of the pose compensation algorithm described in this section is to generate, given a pair of color and depth images, novel corresponding color and depth images depicting a frontal, upright face orientation. Also the center of the face on the input image is aligned with the center of the face in the gallery images of the same person with pixel accuracy.

The proposed technique uses the range image only for face detection and pose estimation and therefore is robust especially under varying pose and illumination conditions, as demonstrated by the experimental results.

The detection of the face in the image is the first step of the algorithm. Segmentation of the head from the body relies on statistical modelling of the head - torso points using a mixture of Gaussians assumption. The parameters of the model are then estimated by means of the Expectation Maximization algorithm and by incorporation of a-priori constraints on the relative dimensions of the body parts, described in detail in [8].

The estimation of 3D head pose, performed next is based on the detection of the nose [8]. After the tip of the nose is localized a 3D line is fitted on the 3D coordinates of pixels on the ridge of the nose. This 3D line defines two of the three degrees of freedom of the face orientation. The third degree of freedom, that is the rotation angle around the nose axis, is then estimated by finding the 3D plane that cuts the face into two bilateral symmetric parts. The error of the above pose estimation algorithm tested on more than 2000 images is less than 2 degrees.

Once the tip of the nose and the pose of the face have been estimated, a 3D coordinate frame aligned with the face is defined centered on the tip of the nose. A warping procedure is subsequently applied on the input depth image to align this local coordinate frame with a reference coordinate frame, which is defined during the training faces using the gallery images, bringing the face in up-right orientation. The transformation between the local and reference coordinate frames is further refined to pixel accuracy by applying the ICP [9] surface registration algorithm between the warped and a reference (gallery) depth image corresponding to claimed person ID.

The rectified depth image contains missing pixel values that are interpolated using a series of steps. Some of the missing values are determined simply by copying corresponding symmetric pixel values from the other side of the face. Remaining missing pixel values are linearly interpolated from neighboring points. The interpolated depth map is subsequently used to rectify the associated color image also using 3D warping (fig. 1).

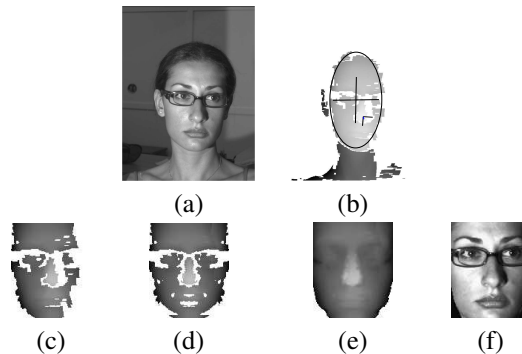


Fig. 1. Pose compensation example. (a) Original color image, (b) original depth image showing detected head blob and estimated local coordinate system fixed on the nose, (c) rectified depth image, (d) symmetry-based interpolation, (e) final linear interpolated depth image, (f) rectified color image.

The proposed pose compensation algorithm is very accurate as will be demonstrated in 5 but also computationally efficient, with total running time is less than 1 sec on a Pentium III 1 Ghz computer.

4. ILLUMINATION COMPENSATION

In this section an algorithm is described that compensates illumination by generating from the input image a novel image relight from a frontal direction. Our approach is inspired by recent work on image-based scene relighting used for rendering realistic images. Image relighting relies on inverting the rendering equation, i.e. the equation that relates the image brightness with the object material and geometry and the illumination of the scene. Given several images of the scene under different conditions this equation may be solved (although an ill-posed problem) to recover the illumination distribution and then use this to re-render the scene under novel illumination.

The first step is therefore to recover the scene illumination from a pair of color and depth images. Assuming that the scene is illuminated by a single light source a technique is adopted that learns the non-linear relationship between the image brightness and light source direction \mathbf{L} using a set of artificially generated bootstrap images.

For each subject in our database we use the reference pose compensated depth image I_r to render N virtual views of the face illuminated from different directions. The set of light source directions is uniformly sampled from a section of the positive hemisphere. To decrease the dimensionality of the problem, from each rendered image a feature vector is extracted containing locally weighted averages of image brightness over M preselected image locations ($M = 30$ in

our experiments). The sample locations are chosen so as to include face areas with similar albedo (i.e. the skin). Feature vectors \mathbf{x}_i , $i = 1, \dots, N$ extracted from all the images, normalized to have zero mean and unit variance, are then used as samples of the M -dimensional illuminant direction function $\mathbf{L} = \mathbf{G}(\mathbf{x})$. An approximation of this function $\tilde{\mathbf{G}}$ using the samples is a regression problem that may be efficiently solved using Support Vector Machines (SVM) [10]. Assume now that we want to compute the similarity between a pose compensated probe image and gallery images of a person j in the gallery. A feature vector \mathbf{x} is computed from the probe image as described previously. Then an estimate of the light source direction is given by $\tilde{\mathbf{G}}_j(\mathbf{x})$ i.e. the SVM regression function computed for the person j during the training phase.

Given the estimate of the light source direction \mathbf{L} relighting the input image with frontal illumination \mathbf{L}_0 is straightforward. Let I_C, I_D be respectively the input pose compensated color and depth images and \tilde{I}_C the illumination compensated image. Then the image irradiance for each pixel \mathbf{u} is approximated by,

$$I_C(\mathbf{u}) = A(\mathbf{u})R(I_D, \mathbf{L}, \mathbf{u}), \quad \tilde{I}_C(\mathbf{u}) = A(\mathbf{u})R(I_D, \mathbf{L}_0, \mathbf{u}) \quad (1)$$

where A is the unknown face albedo or texture function (geometry independent component) and R is a rendering of the surface with constant albedo. Equation 1 is written

$$\tilde{I}_C(\mathbf{u}) = I_C(\mathbf{u}) \frac{R(I_D, \mathbf{L}_0, \mathbf{u})}{R(I_D, \mathbf{L}, \mathbf{u})}$$

i.e. the illumination compensated image is given by multiplication of the input image with a ratio image.




Fig. 2. Illumination compensation example. (a) Original image, (b) $R(I_D, \mathbf{L}, \mathbf{u})$, (c) frontally illuminated image

Figure 2 illustrates the relighting of a side illuminated image.

The same relighting procedure is also applied on training images. Then it is expected that illumination compensated probe and gallery images of the same person will only

differ up to a scale factor since the intensity of the light source may not be recovered. This scale factor is cancelled by taking the logarithm of the images (that makes the factor additive instead of multiplicative) and subsequently subtracting the mean value.

Although the description of above relighting technique considers a single channel image, color images may be handled equally well by applying the same procedure (illuminant estimation and relighting) separately for each channel.

An important advantage of the previously described algorithm is the flexibility in coping with complex illumination conditions by adaptation of the rendering function R above. For example, accounting for attached shadows may be simply achieved by activating shadowing in the rendering engine. On the other hand, from our experience with different rendering models, good results may be also obtained with relatively simple renderings.

5. EXPERIMENTAL RESULTS

The focus of the experimental evaluation was to investigate the improvement achieved by incorporating the proposed pose and illumination schemes into state-of-the-art 2D face recognition algorithms. We have therefore used the Embedded Hidden Markov Model algorithm [11] as a baseline classification algorithm. Two such classifiers are used, one for color images (for practical reason only the red component of the color images was used) and one for depth images. The results of each classifier is a similarity measure for every person in the database. Using the similarity measures associated with color and depth images respectively a combined similarity measure is obtained using the product rule [12].

We have performed several experiments using images of the recorded face database. Training of the classifier was performed using images from the first recording session. On average 3 images per subject depicting different facial expressions were used for training. Testing was performed using all images of the second recording session.

Table 1 demonstrates the recognition rates achieved with the proposed compensation scheme. This is compared with the case that no compensation is performed (the face detection algorithm in [11] was applied in this case), and with manual pose normalization i.e. three points over the eyes and mouth were selected by a human operator and used to rectify the images. Rectification in this case is performed either by 2D affine warping of the images or by 3D warping using depth information as described in section 3. As shown in table 1 the proposed scheme results in significant improvements in the recognition accuracy and its very close to the accuracy achieved by manual image normalization.

Very good results were also obtained by the proposed illumination compensation technique (table 2). We have

	All			Pose		
	C	D	C+D	C	D	C+D
NC	77.4	82.6	84.3	72.2	79.7	81.5
WA	89.7	94.1	96.2	80.1	93.9	95.1
PC	90.8	96.6	98.5	81.5	94.8	95.6
W3D	91.0	96.8	98.9	82.4	95.3	96.2

Table 1. Recognition rates for pose compensation schemes (Color images: C, depth images: D, color + depth: C+D), NC: no compensation, WA: affine warping with manually selected feature points, PC: proposed pose compensation algorithm, W3D: 3D warping with manually selected feature points. The first three columns correspond to all images in the database and the rest three columns correspond to images with large face orientation angle.

	All	Illum.
NC	90.8	90.6
IC	92.7	92.4
CB	87.5	86.4

Table 2. Recognition rates for illumination compensation on color images, NC: no compensation, IC: proposed illumination compensation algorithm, CB: robust similarity measure [13]. The first column corresponds to all images in the database and the second column corresponds to images with strong side illumination.

applied the proposed relighting algorithm to pose compensated images and the resulting images are provided to the 2D color classifier. Also, results obtained with the robust similarity measure proposed in [13] are given for comparison.

In summary we have proposed a new approach for 3D face recognition based on automatic image normalization algorithms exploiting the availability of 3D information. Significant improvements in face classification accuracy were obtained using this scheme. We hope in further improvement of these results in the future using model-based image warping for pose compensation and also by the investigation of efficient reflectance estimation techniques to further enhance illumination compensation.

6. REFERENCES

[1] F. Forster, P. Rummel, M. Lang, and B. Radig, "The hiscore camera: a real time three dimensional and color camera," in *Proc. Int. Conf. Image Processing*, Oct. 2001, vol. 2, pp. 598–601.

[2] G.G. Gordon, "Face recognition based on depth and curvature features," in *Proc. of IEEE Computer So-*

ciety Conf. on Computer Vision and Pattern Recognition, CVPR '92, 1992, pp. 808–810.

- [3] H. T. Tanaka, M. Ikeda, and H. Chiaki, "Curvature-based face surface recognition using spherical correlation. principal directions for curved object recognition," in *Proc. 3rd IEEE Int. Conf. on Automatic Face and Gesture Recognition*, 1998, pp. 372–377.
- [4] C.-S. Chua, F. Han, and Y.-K. Ho, "3d human face recognition using point signature," in *Proc. 4th IEEE Int. Conf. on Automatic Face and Gesture Recognition*, 2000, pp. 233–238.
- [5] Y. Wang, C.-S. Chua, Y.-K. Ho, and Y. Ren, "Integrated 2d and 3d images for face recognition," in *Proc. 11th Int. Conf. on Image Analysis and Processing*, 2001, pp. 48–53.
- [6] S. Tsutsumi, S. Kikuchi, and M. Nakajima, "Face identification using a 3d gray-scale image—a method for lessening restrictions on facial directions," in *Proc. 3rd IEEE Int. Conf. on Automatic Face and Gesture Recognition*, 1998, pp. 306–311.
- [7] K. Chang, K. Bowyer, and P. Flynn, "Face recognition using 2d and 3d facial data," in *Proc. Multimodal User Authentication Workshop*, Santa Barbara, December 2003, to appear.
- [8] S. Malassiotis and M. G. Strintzis, "Real-time head tracking and 3d pose estimation from range data," in *Proc. Int. Conf. Image Processing*, Barcelona, Spain, September 2003.
- [9] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. Pattern Anal. and Mach. Intell.*, vol. 14, no. 2, pp. 239–256, February 1992.
- [10] K.-R. Muller, S. Mika, G. Ratsch, K. Tsuda, and B. Scholkopf, "An introduction to kernel-based learning algorithms," *IEEE Neural Networks*, vol. 12, no. 2, pp. 181–201, May 2001.
- [11] F. Samaria, *Face Recognition Using Hidden Markov Models*, Ph.D. thesis, University of Cambridge, 1984.
- [12] J. Kittler, M. Ballette, J. Czyz, F. Roli, and L. Vanden-dorpe, "Enhancing the performance of personal identity authentication systems by fusion of face verification experts," in *Proc. IEEE Int. Conf. on Multimedia and Expo*, 2002, vol. 2, pp. 581–584.
- [13] H. F. Chen, P. N. Belhumeur, and D. W. Jacobs, "In search of illumination invariants," in *Proc. Conf. Computer Vision and Pattern Recognition*, 2000, pp. I: 254–261.