



Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Thermal and reflectance based personal identification methodology under variable illumination

Ognjen Arandjelović^{a,*}, Riad Hammoud^b, Roberto Cipolla^a

^a Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

^b Delphi Corporation, Delphi Electronics and Safety, Kokomo, IN 46901-9005, USA

ARTICLE INFO

Article history:

Received 17 February 2009

Received in revised form

20 August 2009

Accepted 26 November 2009

Keywords:

Face

Recognition

Thermal

Infrared

Fusion

Illumination

Invariance

ABSTRACT

The objective of this work is to authenticate individuals based on the appearance of their faces. This is a difficult pattern recognition problem because facial appearance is generally greatly affected by the changes in the way a face is illuminated, by the camera viewpoint and partial occlusions, for example due to eye-wear. We describe a fully automatic algorithm that systematically addresses each of these challenges. The main novelty is an algorithm for decision-level fusion of two types of imagery: one acquired in the visual and one acquired in infrared electromagnetic spectrum. Specifically: we examine: (i) the effects of preprocessing of data in each domain, (ii) the fusion of holistic and local facial appearance, and (iii) propose an algorithm for combining the similarity scores in visual and thermal spectra in the presence of prescription glasses and significant pose variations, using a small number of training images (5–7). Our system achieved a high correct identification rate of 97% on a freely available data set containing extreme illumination changes.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

For decades, the personal identification task had shown progress by employing technological means like secret knowledge, such as personal identification numbers, and by using personal possessions, such as identity cards and radio frequency identification chips. As opposed to these means which are generally easy targets for fraud, biometric modalities like facial geometry, ear form and iris are universal and relatively consistent over time. As the cost of biometric sensors, visible and thermal imagers, microphones, and motion sensors continues to sink due to higher demand, biometric systems have the tendency to employ more than a single sensor to identify an individual upon as many as non-redundant biometric modalities. In recent years there have been largely three categories of approaches to face recognition for obtaining complementary information that can be loosely termed multi-sample, multi-sensor and multi-modal. This paper presents a multi-sensory multi-sample face biometric fusion methodology for personal identification.

1.1. Problem challenges and proposed methodology

Variations in head pose and illumination are the most challenging aspects of face recognition. In practice, the effects of

changing pose are usually less problematic and can oftentimes be overcome by acquiring data over a time period e.g. by tracking a face. Consequently, image sequence or image set matching has recently gained a lot of attention in the literature [1–3], as a practical way of obtaining multi-view appearance information. In this paper we too consider the image set matching paradigm. In contrast to pose, illumination changes are much more difficult to deal with for face recognition algorithms: the illumination setup in which a face is imaged is in most cases not possible to control, its physics difficult to accurately model and training data containing typical appearance variability is usually not available.

Thermal spectrum imagery is useful in this regard as it is virtually insensitive to illumination changes, see Fig. 1. On the other hand, it lacks much of the individual, discriminating facial detail contained in visual images. In this sense, the two modalities can be seen as complementing each other. The key idea behind the system presented in this paper is that robustness to extreme illumination changes can be achieved by *fusing* the two. This paradigm will further prove useful when which we consider the difficulty of recognition in the presence of prescription glasses, which are a major problem for recognition using thermal spectrum imagery.

1.2. Paper organization

The remainder of this paper is organized as follows. In the next section we review relevant previous work on face recognition in

* Corresponding author. Tel.: +44 1223 339926.

E-mail addresses: oa214@eng.cam.ac.uk, oa214@cam.ac.uk (O. Arandjelović).

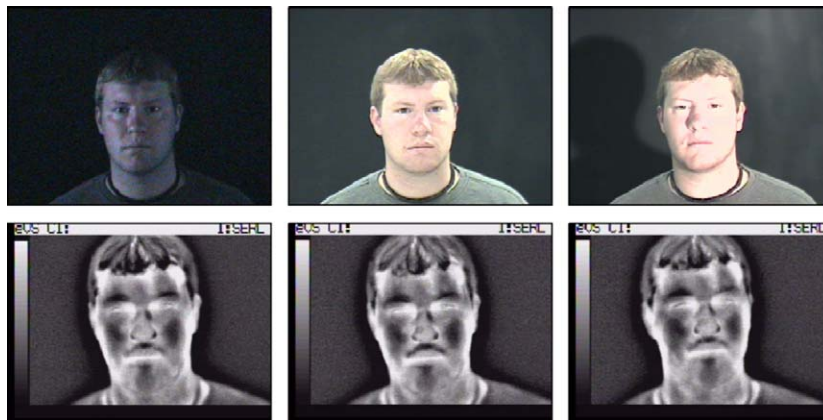


Fig. 1. Illumination changes have a dramatic effects on images acquired in the visible light spectrum (top row). In contrast, thermal imagery (bottom row) shows remarkable invariance.

visual and thermal spectra, as well as multi-sensory techniques. Section 3 describes each of the main components of the proposed system in detail: baseline face set matching is covered first (Section 3.1), followed by sections on data preprocessing and extraction (Section 3.3), recognition using only a single modality (Section 3.4), the proposed modality fusion (Section 3.5) and, finally, dealing with occlusions caused by prescription glasses (Section 3.6). Our empirical evaluation methodology, data sets used and the performance of the proposed algorithm are reported and discussed in Section 4. The paper is concluded with a summary and an outline of promising directions for future research.

2. Previous work

2.1. Mono-sensor based techniques

2.1.1. Optical sensors

By far the most widely used type of sensors in face recognition systems are optical sensors, as a result of their high availability and low cost. An optical camera records the energy of incident light in the visible spectrum reflected from the face surface. Thus, it implicitly captures the variation in the surface shape and reflectance. However, it is also dependent on the lighting conditions in which the images are acquired. In order to overcome this problem, a larger number of face recognition algorithms has been proposed which can be broadly into two well-known categories: statistical appearance-based and model-based methods [4].

Appearance-based methods, as their name suggests, perform recognition directly from the way faces appear in images, interpreting them as ordered collections of pixels. Faces are typically represented as vectors in the image space; discrimination between individuals is then performed by employing statistical models that explain inter- and/or intra- personal appearance changes.

The main limitation of purely appearance-based recognition is that of limited generalization ability, especially in the presence of greatly varying illumination conditions. On the other hand, an appealing property of this group of algorithms is that they can be typically extended to recognize using sets or sequences of faces. Eigenfaces, for example, can be used on a per-image basis, with recognition decision cast using majority vote [5]. A similar voting approach was also successfully used with local features in [6], which were extracted by tracking a face using a Gabor Wavelet Network [6–8]. In [9] video information is used only in the

training stage to construct person-specific PCA spaces, *self-eigenfaces*, while verification was performed from a single image using the Distance from Feature Space criterion. Classifiers using different eigenfeature spaces were used in [10] and combined using the sum rule [11]. Better use of training data is made with various discriminative methods such as Fisherfaces, which can be used to estimate *database-specific* optimal projection [12]. An interesting extension of appearance correlation-based recognition to matching sets of faces was proposed by Yamaguchi et al. [13]. The so-called mutual subspace method (MSM) has since gained considerable attention in the literature. In MSM, linear subspaces describing appearance variations within sets or sequences are matched using canonical correlations [14,15].

However, all of the aforementioned appearance-based methods poorly generalize in the presence of large illumination changes. Thus, model-based methods have been proposed to deal with this problem. These formulate models of image formation with the intention of recovering: (i) mainly person-specific (e.g. face albedo or shape) and (ii) extrinsic, nuisance variables (e.g. illumination direction, or head yaw). The key challenge lies in coming up with models for which the parameter estimation problem is not ambiguous or ill-conditioned.

The active appearance model [16], for example, was proposed for describing objects that vary both in shape and appearance. A deformable, triangular mesh is fitted to an image of a face, guided by combined statistical models of shape and shape-free appearance, so as to best explain the observed image. The 3D Morphable Model can be seen as a 3D extension of this approach. This model consists of albedo values at the nodes of a 3-dimensional triangular mesh describing face geometry. Model fitting is performed by combining a Gaussian prior on the shape and texture of human faces with photometric information from an image [17].

An attractive feature of model-based approaches is that they explicitly model both intrinsic and extrinsic variables. On the other hand, they suffer from convergence problems in the presence of background clutter or facial occlusions such as glasses, unlike the method proposed in this paper. Furthermore and importantly for this work they often require high quality image input and struggle with non-Lambertian effects or multiple light sources.

2.1.2. Thermal infrared-sensor based techniques

A number of studies suggests that face recognition in the thermal spectrum offers some distinct advantages over that in the visible spectrum, including invariance to ambient illumination

changes [18–21]. This is due to the fact that a thermal infrared sensor measures the energy of heat radiation emitted by the face rather than the reflected light. In outdoor environments, and particularly in direct sunlight, illumination invariance only holds true to good approximation for the long-wave infrared (LWIR: 8–14 μm) spectrum, which is fortunately measured by the less expensive uncooled thermal infrared camera technology. Human skin has high emissivity in the medium-wave infrared (MWIR: 3–5 μm) spectrum and even higher emissivity in the LWIR spectrum making face imagery by and large invariant to illumination variations in these spectra.

Appearance-based face recognition algorithms applied to thermal infrared imaging consistently performed better than when applied to visible imagery under variable lighting conditions [22,23,19,24]. Further performance improvements can be achieved using decision-based fusion [19]. In contrast to other techniques, Srivastana and Liu [25], performed face recognition in the space of Bessel function parameters. First, they decompose each infrared face image using Gabor filters. Then, they represent the face by modelling the marginal density of the Gabor filter coefficients using Bessel functions. This approach has further been improved by Buddhharaju et al. [26]. Recently, Friedrich and Yeshurun [27] showed that IR-based recognition is also less sensitive to changes in head pose and facial expression.

A thermal sensor generates imaging features that uncover thermal characteristics of the face pattern. Another advantage of thermal infrared imaging in face recognition is the existence of a direct relationship to underlying physical anatomy such as vasculature. Indeed, thermal face recognition algorithms attempt to take advantage of such anatomical information of the human face as unique signatures. The use of vessel structure for human identification has been studied during recent years using traits such as hand vessel patterns [28,29], finger vessel patterns [30,31] and vascular networks from thermal facial images [32]. In [33] a novel methodology that consists of a statistical face segmentation and a physiological feature extraction algorithm, and a matching procedure of the vascular network from thermal facial imagery has been proposed.

The down side of employing near infrared and thermal infrared sensors is that glare reflections and opaque regions appear in presence of subjects wearing prescription glasses, plastic and sun glasses. For a large proportion of individuals the regions around the eyes—that is an area of high interest to face recognition systems—become occluded and therefore less discriminant [34,35].

2.2. Multi-sensor based techniques

In the biometric literature several classifiers have been used to concatenate and consolidate the match scores of multiple independent matchers of biometric traits [36–40]. In [41] a HyperBF network is used to combine matchers based on voice and face features. Ross and Jain [42] use decision tree and linear discriminant classifiers for classifying the match scores pertaining to the face, fingerprint and hand geometry modalities. In [43] three different color channels of a face image are independently subjected to LDA and then combined.

Recently, several successful attempts have been made to fuse the visual and thermal infrared modalities to increase the performance of face recognition [44,45,21,46,47,22,41,42,48,49]. Visible and thermal sensors are well-matched candidates for image fusion as limitations of imaging in one spectrum seem to be precisely the strengths of imaging in the other. Indeed, as the surface of the face and its temperature have nothing in common,

it would be beneficial to extract and fuse cues from both sensors that are not redundant and yet complementary.

In [44] two types of visible and thermal fusion technique have been proposed. The first fuses low-level data while the second fuses matching distance scores. Data fusion was implemented by applying pixel-based weighted averaging of co-registered visual and thermal images. Decision fusion was implemented by combining the matching scores of individual recognition modules.

The fusion at the score level is the most commonly considered approach in the biometric literature [50]. Cappelli et al. [51] use a *double sigmoid function* for score normalization in a multi-biometric system that combines different fingerprint matchers. Once the match scores output by multiple matchers are transformed into a common domain they can be combined using simple fusion operators such as the sum of scores, product of scores or order statistics (e.g., maximum/minimum of scores or median score). Our proposed method falls into this category of multi-sensor fusion at the score level. To deal with occlusions caused by eyeglasses in thermal imagery, Heo et al. [44] used a simple ellipse fitting technique to detect the circle-like eyeglass regions in the IR image and replaced them with an average eye template. Using a commercial face recognition system, Facelt [52], they demonstrated improvements in face recognition accuracy. Our method differs both in the glasses detection stage, which uses a principled statistical model of appearance variation, and in the manner it handles detected occlusions. Instead of using the average eye template, which carries no discriminative information, we segment out the eye region from the infrared data, effectively placing more weight on the discriminative power of the same region extracted from the filtered, visual imagery.

3. Method details

In the sections that follow we explain our system in detail, the main components of which are conceptually depicted in Fig. 2.

3.1. Matching image sets

In this paper we are dealing with face recognition from sets of images, both in the visual and thermal spectrum. We will demonstrate that illumination invariance and recognition robustness in the presence of prescription glasses can be achieved using a combination of simple data preprocessing (Section 3.3), a combination of holistic and local features (Section 3.4) and the fusion of the two modalities (see Section 3.5). These stages normalize for the bulk of appearance changes caused by extrinsic (non-person-specific) factors. Hence, the requirements for our basic set-matching algorithm are those of: (i) some pose generalization and (ii) robustness to noise. We will show that the two criteria are successfully met by modelling appearance variations within an image set using a low-dimensional linear subspace and then by comparing two such subspaces by finding the most similar modes of variation within them.

The appearance modelling step is a simple application of principal component analysis (PCA) without mean subtraction. Formally, given a data matrix $\mathbf{d} \in \mathbb{R}^{D \times N}$ (each column representing a rasterized image), the corresponding subspace is spanned by the eigenvectors of the matrix $\mathbf{C} = \mathbf{d}\mathbf{d}^T$ corresponding to the largest eigenvalues. We used 5D subspaces, as sufficiently expressive to on average explain over 90% of data variation within intrinsically low-dimensional face appearance changes within a set.

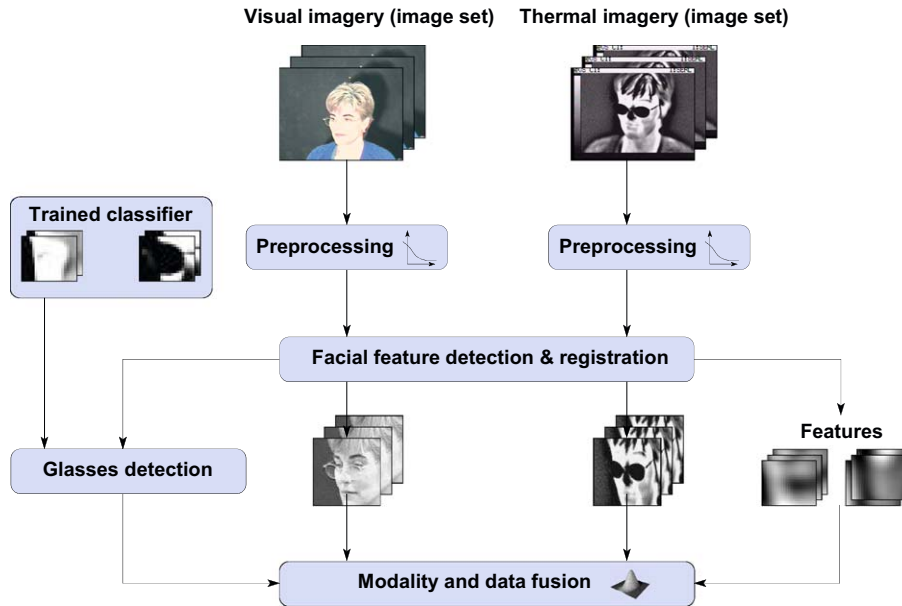


Fig. 2. Our system consists of three main modules performing (i) data preprocessing and registration, (ii) glasses detection and (iii) fusion of holistic and local face representations using visual and thermal modalities.

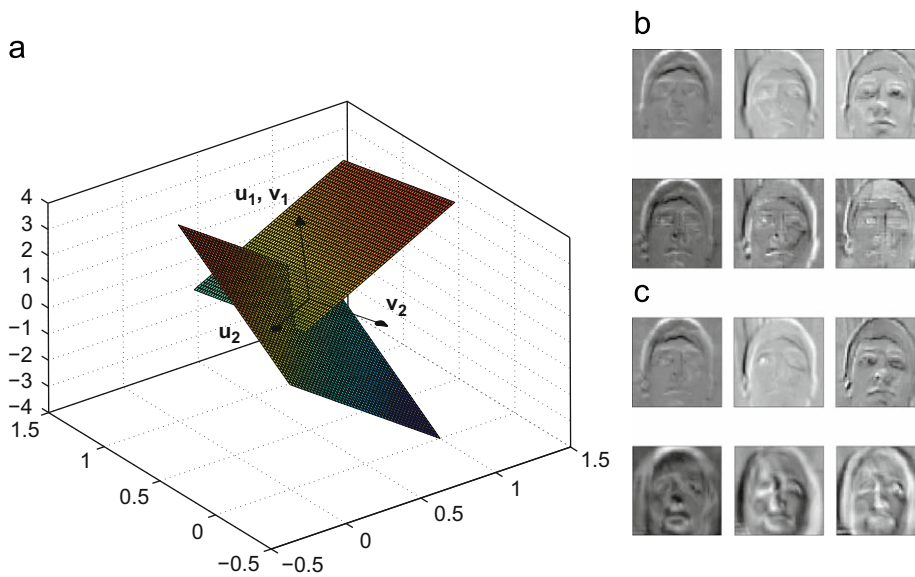


Fig. 3. An illustration of the concept of principal angles and principal vectors in the case of two 2D subspaces embedded in a 3D space. As two such subspaces necessarily intersect, the first pair of principal vectors is the same (i.e. $\mathbf{u}_1 = \mathbf{v}_1$). However, the second pair is not, and in this case forms the second principal angle of $\cos^{-1} \rho_2 = \cos^{-1}(0.8084) \approx 36^\circ$. The top three pairs of principal vectors, displayed as images, when the subspaces correspond to image sets of the same and different individuals are displayed in (b) and (c) (top rows corresponds to \mathbf{u}_i , bottom rows to \mathbf{v}_i). In (b), the most similar modes of pattern variation, represented by principal vectors, are very much alike in spite of different illumination conditions used in data acquisition.

3.2. Principal angles

Principal, or canonical, angles $0 \leq \theta_1 \leq \dots \leq \theta_D \leq (\pi/2)$ between two D -dimensional linear subspaces U_1 and U_2 are recursively uniquely defined as the minimal angles between any two vectors of the subspaces:

$$\rho_i = \cos \theta_i = \max_{\mathbf{u}_i \in U_1} \max_{\mathbf{v}_i \in U_2} \mathbf{u}_i^T \mathbf{v}_i \quad (1)$$

subject to the orthonormality condition:

$$\mathbf{u}_i^T \mathbf{u}_i = \mathbf{v}_i^T \mathbf{v}_i = 1, \quad \mathbf{u}_i^T \mathbf{u}_j = \mathbf{v}_i^T \mathbf{v}_j = 0, \quad j = 1, \dots, i-1. \quad (2)$$

We will refer to \mathbf{u}_i and \mathbf{v}_i as the i -th pair of *principal vectors*, see Fig. 3(a). The quantity ρ_i is also known as the i -th canonical correlation [15]. Intuitively, the first pair of principal vectors corresponds to the most similar modes of variation within two linear subspaces; every next pair to the most similar modes orthogonal to all previous ones. We quantify the similarity of subspaces U_1 and U_2 , corresponding to two face sets, by the cosine of the smallest angle between two vectors confined to them i.e. ρ_1 .

This interpretation of principal vectors motivates the suitability of canonical correlations as a similarity measure when subspaces U_1 and U_2 correspond to face images. First, the



Fig. 4. Shown is the original image in the visual spectrum with detected facial features marked by yellow circles (left), the result of affine warping the image to the canonical frame (centre) and the final registered and cropped facial image. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

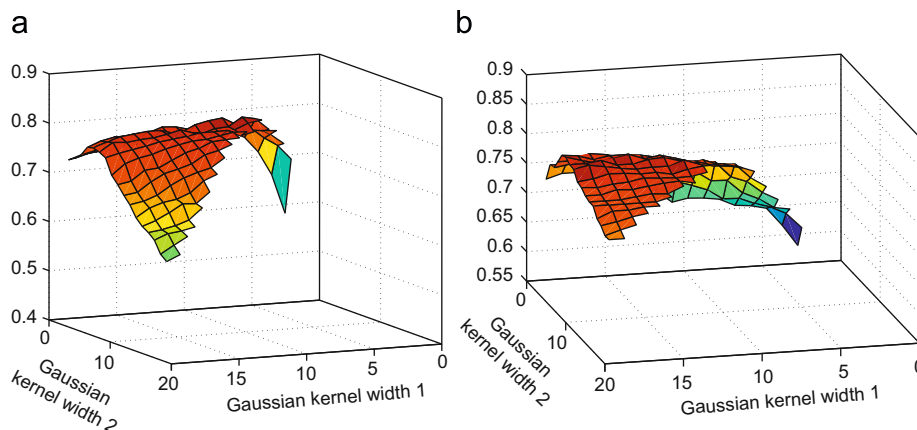


Fig. 5. The optimal combination of the lower and upper band-pass filter thresholds is estimated from a small training corpus. The plots show the recognition rate using a single modality, (a) visual and (b) thermal, as a function of the widths W_1 and W_2 of the two Gaussian kernels in (4). It is interesting to note that the optimal band-pass filter for the visual spectrum passes a rather narrow, mid-frequency band, whereas the optimal filter for the thermal spectrum is in fact a *low-pass* filter.

empirical observation that face appearance varies smoothly as a function of camera viewpoint [1,53] is implicitly exploited: since the computation of the most similar modes of appearance variation between sets can be seen as an efficient “search” over entire subspaces, *generalization* by means of linear pose interpolation and extrapolation is inherently achieved. This concept is further illustrated in Fig. 3(b and c). Furthermore, by being dependent on only a single (linear) direction within a subspace, by employing the proposed similarity measure the bulk of data in each set, deemed not useful in a specific set-to-set comparison, is thrown away. In this manner *robustness to missing data* is achieved.

An additional appealing feature of comparing two subspaces in this manner is contained in its computational efficiency. If \mathbf{B}_1 and \mathbf{B}_2 are orthonormal basis matrices corresponding to U_1 and U_2 , then writing the Singular Value Decomposition (SVD) of the matrix $\mathbf{B}_1^T \mathbf{B}_2$:

$$\mathbf{M} = \mathbf{B}_1^T \mathbf{B}_2 = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T. \quad (3)$$

The i -th canonical correlation ρ_i is then given by the i -th singular value of \mathbf{M} i.e. $\Sigma_{i,i}$, and the i -th pair of principal vectors \mathbf{u}_i and \mathbf{v}_i by, respectively, $\mathbf{B}_1 \mathbf{U}$ and $\mathbf{B}_2 \mathbf{V}$ [54]. Seeing that in our case \mathbf{M} is a 5×5 matrix and that we only use the largest canonical correlation, ρ_1 can be rapidly computed as the largest eigenvalue of $\mathbf{M} \mathbf{M}^T$ [55].

3.3. Data preprocessing and feature extraction

The first stage of our system involves coarse normalization of pose and illumination. Pose changes are accounted for by in-plane

registration of images, which are then passed through quasi-illumination-invariant image filters.

We register all faces, both in the visual and thermal domain, to have the salient facial features aligned. Specifically, we propose to explicitly align the eyes and the mouth. This choice of features is motivated by their characteristic appearance in both thermal and visual spectra, which makes automatic detection very reliable (e.g. see [56–60]). In this work, we used the method proposed by Arandjelović and Zisserman [56]. The 3 point correspondences, between the detected and the canonical features’ locations, uniquely define an affine transformation which is applied to the original image. Finally, faces are then cropped to 80×80 pixels, as shown in Fig. 4.

Coarse brightness normalization is performed by band-pass filtering the images [56,61]. The aim is to reduce the amount of high-frequency noise as well as extrinsic appearance variations confined to a low-frequency band containing little discriminating information. Most obviously, in visual imagery, the latter are caused by illumination changes, owing to the smoothness of the surface and albedo of faces [62].

We consider the following type of a spatial-domain band-pass filter:

$$\mathbf{I}_F = \mathbf{I} * \mathbf{G}_{\sigma=W_1} - \mathbf{I} * \mathbf{G}_{\sigma=W_2}, \quad (4)$$

which has two parameters—the widths W_1 and W_2 of isotropic Gaussian kernels. The optimal values of these are estimated from a small training corpus of individuals in varying illumination. Fig. 5 shows the recognition rate across the corpus as the values of the two parameters are varied. The optimal values were found to be 2.3 and 6.2 for visual data; the optimal filter for thermal data was found to be a *low-pass* filter with $W_2 = 2.8$ (i.e. W_1 was found to be very large). This is indeed expected, as thermal imagery is not affected by any predominantly low-frequency extrinsic

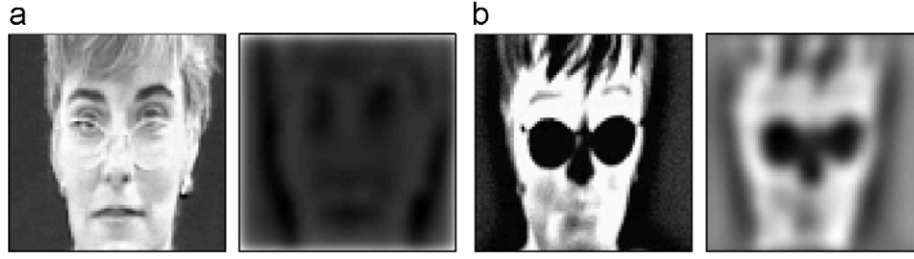


Fig. 6. Effects of the optimal band-pass filters on registered and cropped faces in (a) visual and (b) thermal spectra.

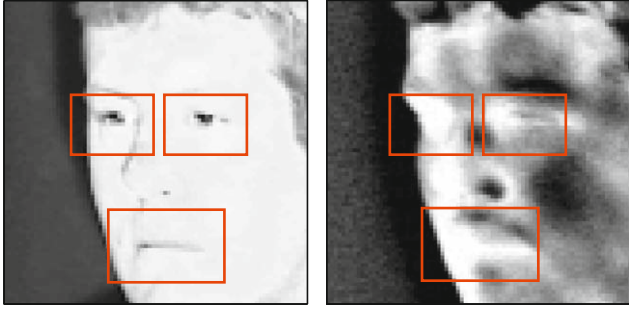


Fig. 7. In both the visual and the thermal spectrum our algorithm combines the similarities obtained by matching the holistic face appearance and the appearance of three salient local features—the eyes and the mouth.

factors, but can be seen to contain a rather large amount of high frequency noise.

Examples of optimally filtered, registered and cropped images are shown in Fig. 6. It is important to note from Fig. 5 that the recognition rate varied smoothly with changes in kernel widths, showing that the method is not very sensitive to their exact values, which is suggestive of good generalization to unseen data.

Finally, we also propose to further scale the result of filtering of visual data by the smooth version of its original image:

$$\hat{\mathbf{I}}_F(x, y) = \mathbf{I}_F(x, y) ./ (\mathbf{I} * \mathbf{G}_{\sigma = w_2}), \quad (5)$$

where $./$ represents element-wise division. The purpose of local scaling is to equalize edge strengths in dark (weak edges) and bright (strong edges) regions of the face; this is similar to the self-quotient image of Wang et al. [63]. This step further improves the robustness of the representation to illumination changes, as we demonstrate in Section 4.

3.4. Single modality-based recognition

We compute the similarity of two individuals using only a single modality (visual or thermal) by combining the holistic face representation described in Section 3.3 and a representation based on local image patches. The idea behind using patches is to gain further viewpoint robustness. Since smaller surface areas of the face are closer to planar than the entire face, when affine registered they exhibit a greater degree of appearance invariance. This is supported by previous empirical accounts in the literature, e.g. see [3,64,65].

As before, we use the eyes and the mouth as the most discriminative regions by extracting rectangular patches centred at the detections, see Fig. 7. We treat holistic and local patch-based similarities as independent, the overall similarity score thus being obtained by weighted summation:

$$\rho_{v/t} = \omega_h \cdot \rho_h + \underbrace{\omega_m \cdot \rho_m}_{\text{Local patches contribution}} + (1 - \omega_h - \omega_m) \cdot \rho_e, \quad (6)$$

where ρ_m , ρ_e and ρ_h are the scores of separately matching, respectively, the mouth, the eyes and the entire face regions, and ω_h and ω_m the weighting constants.

The optimal values of the weights were estimated from the offline training corpus. As expected, eyes were shown to carry a significant amount of discriminative information, as for the visual spectrum we obtained $\omega_e = 0.3$. On the other hand, the mouth region, highly variable in appearance in the presence of facial expression changes, was found not to improve recognition (i.e. $\omega_m \approx 0.0$).

The relative magnitudes of the weights were found to be different in the thermal spectrum, both the eye and the mouth region contributing equally to the overall score: $\omega_m = 0.1$, $\omega_h = 0.8$. Notice the rather insignificant contribution of individual facial features. This is most likely due to inherently spatially slowly varying nature of heat radiated by the human body.

3.5. Fusing modalities

Until now we have focused on deriving a similarity score between two individuals given sets of images in either thermal or visual spectrum. A combination of holistic and local features was employed in the computation of both. However, the greatest power of our system comes from the fusion of the two modalities.

Given ρ_v and ρ_t , the similarity scores corresponding to visual and thermal data, we compute the joint similarity as:

$$\rho_f = \omega_v(\rho_v) \cdot \rho_v + [1 - \omega_v(\rho_v)] \cdot \rho_t. \quad (7)$$

Notice that the weighting factors are no longer constants, but *functions*. The key idea is that if the visual spectrum match is very good (i.e. ρ_v is close to 1.0), we can be confident that illumination difference between the two images sets compared is mild and well compensated for by the visual spectrum preprocessing of Section 3.3. In this case, visual spectrum should be given relatively more weight than when the match is bad and the illumination change is likely more drastic. The value of $\omega_v(\rho_v)$ can then be interpreted as statistically the optimal choice of the mixing coefficient ω_v given the visual domain similarity ρ_v . Formalizing this we can write

$$\omega_v(\rho_v) = \underset{\omega}{\operatorname{argmax}} p(\omega | \rho_v), \quad (8)$$

or, equivalently

$$\omega_v(\rho_v) = \underset{\omega}{\operatorname{argmax}} \frac{p(\omega, \rho_v)}{p(\mu)}. \quad (9)$$

Under the assumption of a uniform prior on the confusion margin, $p(\mu)$

$$p(\alpha | \mu) \propto p(\alpha, \rho_v), \quad (10)$$

and

$$\omega_v(\rho_v) = \underset{\omega}{\operatorname{argmax}} p(\omega, \rho_v). \quad (11)$$

3.5.1. Learning the weighting function

The function $\omega_v \equiv \omega_v(\rho_v)$ is estimated in three stages: first (i) we estimate $p(\omega_v, \rho_v)$, then (ii) compute $\omega(\rho_v)$ using (11) and finally (iii) make an analytic fit to the obtained marginal distribution. Step (i) is challenging and we describe it next.

Iterative density estimate: The principal difficulty of estimating $p(\omega_v, \rho_v)$ is of practical nature: in order to obtain an accurate estimate (i.e. a well-sampled distribution), a prohibitively large training database is needed. Instead, we employ a heuristic alternative. Much like before, the estimation is performed using the offline training corpus.

Our algorithm is based on an iterative incremental update of the density, initialized as uniform over the domain $\omega_v, \rho_v \in [0, 1]$. We iteratively simulate matching of an unknown person against a set gallery individuals. In each iteration of the algorithm, these are randomly drawn from the offline training database. Since the ground truth identities of all persons in the offline database is known, for each $\omega_v = k\Delta\omega_v$ we can compute (i) the initial visual spectrum similarity $\rho_v^{p,p}$ of the novel and the corresponding gallery sequences, and (ii) the resulting separation $\delta(k\Delta\omega_v)$ i.e. the difference between the similarities of the test set and the set corresponding to it in identity, and that between the test set and the most similar set that does *not* correspond to it in identity. This gives us information about the usefulness of a particular value of ω_v for observed $\rho_v^{p,p}$. Hence, the density estimate $\hat{p}(\omega_v, \rho_v)$ is then updated at $(k\Delta\omega_v, \rho_v^{p,p})$, $k = 1 \dots$. We increment it proportionally to $\delta(k\Delta\omega_v)$ after passing through a y -axis shifted sigmoid function:

$$\hat{p}(k\Delta\omega_v, \rho_v^{p,p})_{[n+1]} = \hat{p}(k\Delta\omega_v, \rho_v^{p,p})_{[n]} + [\text{sig}(C \cdot \delta(k\Delta\omega_v)) - 0.5], \quad (12)$$

where subscript $[n]$ signifies the n -th iteration step and

$$\text{sig}(x) = \frac{1}{1 + e^{-x}}, \quad (13)$$

as shown in Fig. 8(a). The sigmoid function has the effect of reducing the overly confident weight updates for the values of ω_v that result in extremely good or bad separations $\delta(k\Delta\omega_v)$. The purpose of this can be seen by noting that we are using separation as a proxy for the statistical goodness of ω_v , while in fact attempting to maximize the average recognition rate (i.e. the average number of cases for which $\delta(k\Delta\omega_v) > 0$). This is one of the main differences of the described algorithm to that of Arandjelović and Cipolla proposed in [66].

Fig. 9 summarizes the proposed offline learning algorithm. An analytic fit to $\omega_v(\rho_v)$ in the form $(1 + e^a)/(1 + e^{a/\rho_v})$ is shown in Fig. 8(b).

Input: visual data $d_v(\text{person}, \text{illumination})$,
thermal data $d_t(\text{person}, \text{illumination})$.
Output: density estimate $\hat{p}(\omega_v, \rho_v)$.

1: Initialization

$$\hat{p}(\omega_v, \rho_v) = 0,$$

2: Iteration

for all illuminations i, j and persons p

3: Iteration

for all $k = 0, \dots, 1/\Delta\omega_v$, $\omega_v = k\Delta\omega_v$

5: Separation given ω

$$\delta(k\Delta\omega_v) = \min_{q \neq p} [\omega_v \rho_v^{p,p} + (1 - \omega_v) \rho_t^{p,p} - \omega_v \rho_v^{p,q} + (1 - \omega_v) \rho_t^{p,q}]$$

6: Update density estimate

$$\hat{p}(k\Delta\omega_v, \rho_v^{p,p}) = \hat{p}(k\Delta\omega_v, \rho_v^{p,p}) + \left[\text{sig}(C \cdot \delta(k\Delta\omega_v)) - 0.5 \right]$$

7: Smooth the output

$$\hat{p}(\omega_v, \rho_v) = \hat{p}(\omega_v, \rho_v) * \mathbf{G}_{\sigma=0.05}$$

8: Normalize to unit integral

$$\hat{p}(\omega_v, \rho_v) = \hat{p}(\omega_v, \rho_v) / \int_{\omega_v} \int_{\rho_v} \hat{p}(\omega_v, \rho_v) d\rho_v d\omega_v$$

Fig. 9. The proposed fusion learning algorithm, applied offline.

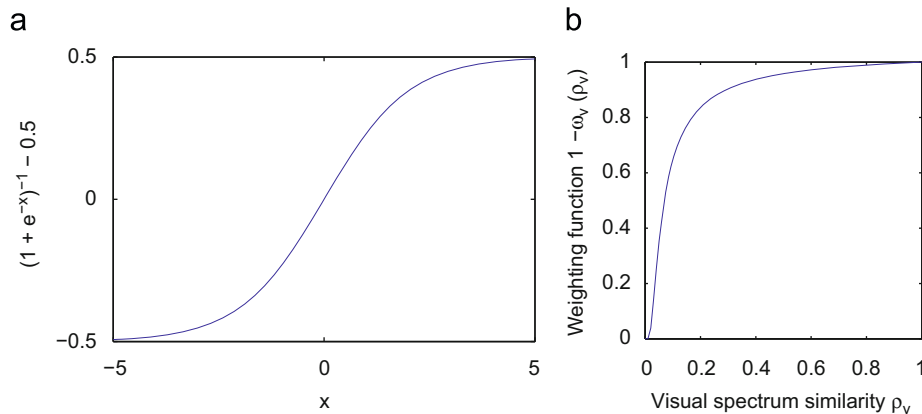


Fig. 8. The contribution of visual matching, as a function of the similarity of visual imagery. A low similarity score between image sets in the visual domain is indicative of large illumination changes and consequently our algorithm leans that more weight should be placed on the illumination-invariant thermal spectrum: (a) y -axis shifted sigmoid function and (b) weighting function.

3.6. Prescription glasses

The appeal of using the thermal spectrum for face recognition stems mainly from its invariance to illumination changes, in sharp contrast to visual spectrum data. The exact opposite is true in the case of prescription glasses, which appear as dark patches in thermal imagery, see Fig. 6. The practical importance of achieving recognition robustness in the presence of corrective eyewear can be seen by noting that in the US in 2000 roughly 96 million people, or 34% of the total population, wore prescription glasses [67].

In our system, the otherwise undesired, gross appearance distortion that glasses cause in thermal imagery is used to help recognition by detecting their presence. If the subject is not wearing glasses, then both holistic and all local patches-based face representations can be used in recognition; otherwise the eye regions in thermal images are ignored as they contain no useful recognition (discriminative) information.

3.6.1. Glasses detection

We detect the presence of glasses by building representations for the left eye region (due to the symmetry of faces, a detector for only one side is needed) with and without glasses, in the thermal

spectrum. The foundations of our classifier are laid out in Section 3.1. Appearance variations of the eye region with out without glasses are represented by two 6D linear subspaces estimated from the training data corpus, see Fig. 10 for examples of training data used for subspace estimations. The linear subspace corresponding to eye region patches extracted from a set of thermal imagery of a novel person is then compared with “glasses on” and “glasses off” subspaces using principal angles. The presence of glasses is deduced when the corresponding subspace results in a higher similarity score. We obtain close to flawless performance on our data set (also see Section 4 for description), as shown in Fig. 11(a,b). Good discriminative ability of principal angles in this case is also supported by visual inspection of the “glasses on” and “glasses off” subspaces; this is illustrated in Fig. 11(c) which shows the first two dominant modes of each, embedded in the 3D principal subspace.

The presence of glasses severely limits what can be achieved with thermal imagery, the occlusion heavily affecting both the holistic face appearance as well as that of the eye regions. This is the point at which our method heavily relies on decision fusion with visual data, limiting the contribution of the thermal spectrum to matching using mouth appearance only i.e. setting $\omega_m = 1.0$ in (6).

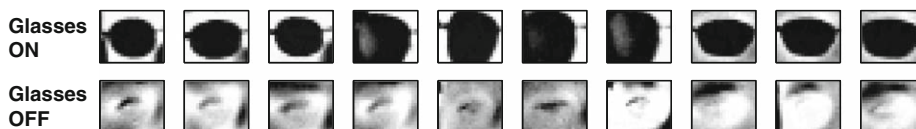


Fig. 10. Shown are examples of glasses-on (top) and glasses-off (bottom) thermal data used to construct the corresponding appearance models for our glasses detector.

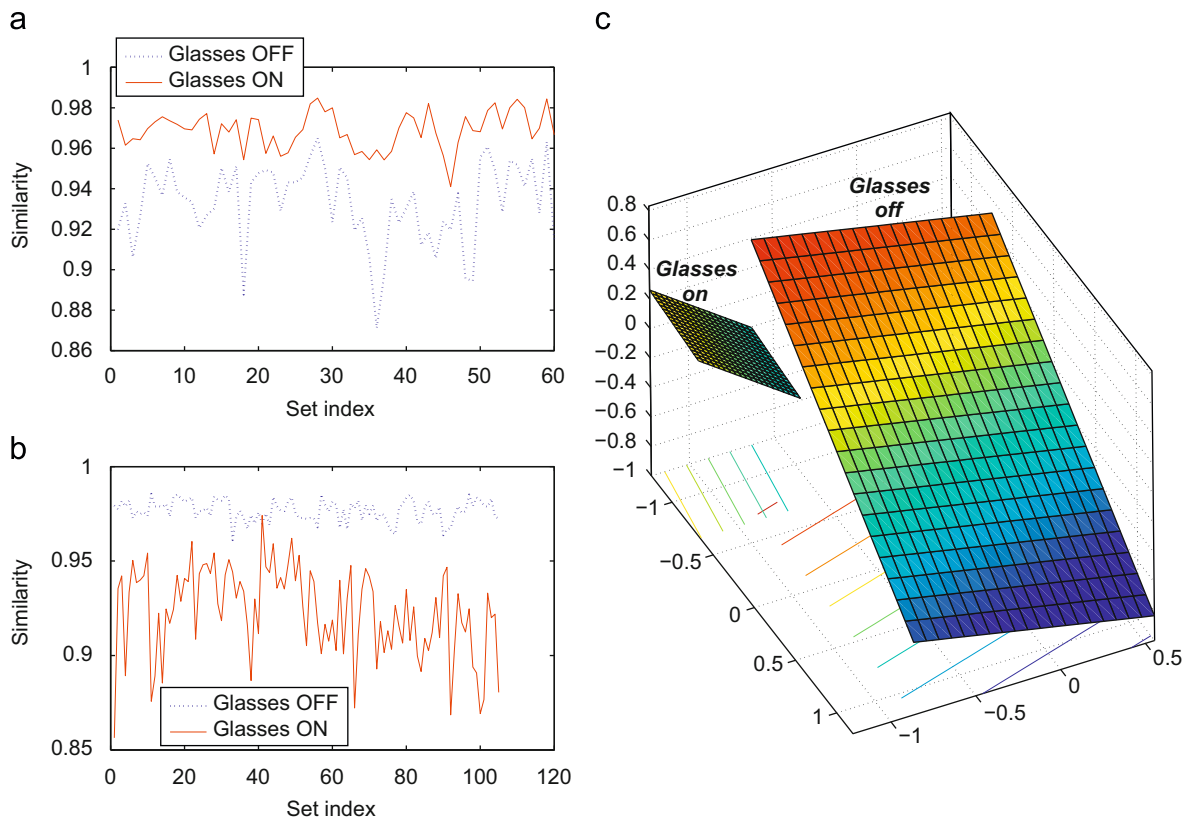


Fig. 11. (a) Inter- and (b) intra-class (glasses on and off) similarities across our data set. (c) Good discrimination by principal angles is also motivated qualitatively as the subspaces modelling appearance variations of the eye region with and without glasses on show very different orientations even when projected to the 3D principal subspace. As expected, the “glasses off” subspace describes more appearance variation, as illustrated by the proportionally greater extent of the linear patch representing it in the plot. (a) Glasses ON, (b) glasses OFF and (c) model subspaces.



Fig. 12. Each row corresponds to an example of a single training (or test) set of images used for our algorithm in (a) the visual and (b) the thermal spectrum. Note the extreme changes in illumination, as well as that in some sets the user is wearing glasses and in some not. (a) Visual and (b) thermal.

4. Empirical evaluation

We evaluated the described system on the “Dataset 02: IRIS Thermal/Visible Face Database” subset of the *Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS)* database,¹ freely available for download at <http://www.cse.ohio-state.edu/OTCBVS-BENCH/>. this database contains 30 individuals, 11 roughly matching poses in visual and thermal spectra and large illumination variations (some of these are exemplified in Fig. 12). Images were acquired using the Raytheon Palm-IR-Pro camera in the thermal and Panasonic WV-CP234 camera in the visual spectrum, in the resolution of 240×320 pixels.

¹ IEEE OTCBVS WS Series Bench; DOE University Research Program in Robotics under Grant DOE-DE-FG02-86NE37968; DOD/TACOM/NAC/ARC Program under Grant R01-1344-18; FAA/NSSA Grant R01-1344-48/49; Office of Naval Research under Grant #N000143010022.

Our algorithm was trained using all images in a single illumination in which all three salient facial features could be detected. This typically resulted in 7–8 images in the visual and 6–7 in the thermal spectrum, see Fig. 13, and roughly $\pm 45^\circ$ yaw range, as measured from the frontal face orientation.

The performance of the algorithm was evaluated both in 1-to- N and 1-to-1 matching scenarios. In the former case, we assumed that test data corresponded to one of people in the training set and recognition was performed by associating it with the closest match. Verification (or 1-to-1 matching, “is this the same person?”) performance was quantified by looking at the true positive admittance rate for a threshold that corresponds to 1 admitted intruder in 100.

4.1. Results

A summary of 1-to- N matching results is shown in Table 1. Firstly, note the poor performance achieved using both raw visual

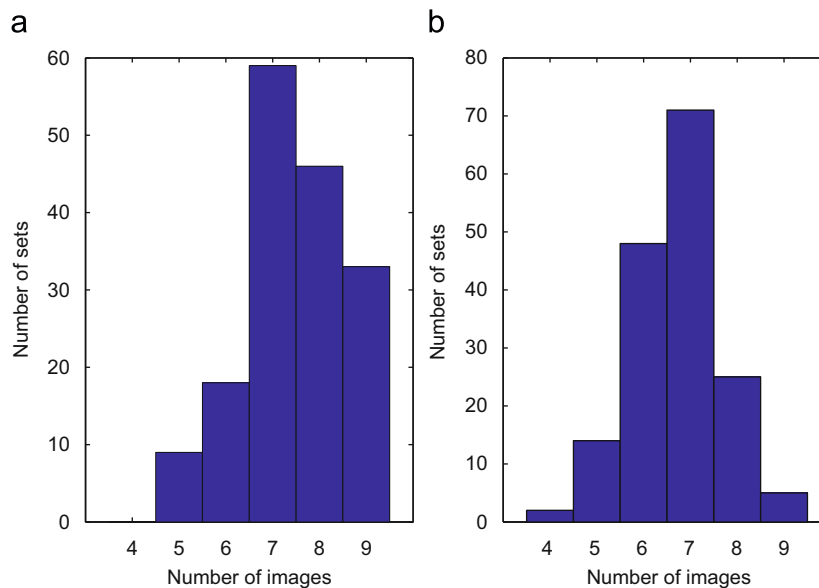


Fig. 13. Shown are histograms of the number of images per person used to train our algorithm. Depending on the exact head poses assumed by the user we typically obtained 7–8 visual spectrum images and typically a slightly lower number for the thermal spectrum. The range of yaw angles covered is roughly $\pm 45^\circ$ measured from the frontal face orientation. (a) Visual and (b) thermal.

Table 1

Shown is the average rank-1 recognition rate using different representations across all combinations of illuminations.

Representation	Recognition
Visual	
Holistic raw data	0.58
Holistic, band-pass	0.78
Holistic, SQI filter	0.85
Local+holistic data fusion, SQI filter	0.87
Thermal	
Holistic raw data	0.74
Holistic raw w/ glasses detection	0.77
Holistic, low-pass filter	0.80
Local+holistic data fusion, low-pass filter	0.82
Multi-sensory fusion	
w/o glasses detection	0.90
w/ glasses detection	0.97

Note the performance increase with each of the main features of our system: image filtering, combination of holistic and local features, modality fusion and prescription glasses detection.

as well as raw thermal data. The former is suggestive of challenging illumination changes present in the OTCBVS data set (Fig. 14). This is further confirmed by significant improvements gained with both band-pass filtering and the self-quotient image which increased the average recognition rate for, respectively, 35% and 47%. The same is corroborated by the receiver-operator characteristic curves in Fig. 15 and 1-to-1 matching results in Table 2.

On the other hand, the reason for low recognition rate of raw thermal imagery is twofold: it was previously argued that the two main limitations of this modality are the inherently lower discriminative power and occlusions caused by prescription glasses. The addition of the glasses detection module is of little help at this point—some benefit is gained by steering away from misleadingly good matches between any two people wearing glasses, but it is limited in extent as a very discriminative region

of the face is lost. Furthermore, the improvement achieved by optimal band-pass filtering in thermal imagery is much more modest than with visual data, increasing performance, respectively, by 35% and 8%. Similar increase was obtained in true admittance rate (42% vs. 8%), see Table 2.

Neither the eyes or the mouth regions, in either the visual or thermal spectrum, proved very discriminative when used in isolation, see Fig. 14. Only 10–12% true positive admittance was achieved, as shown in Table 3. However, the proposed fusion of holistic and local appearance offered a consistent and statistically significant improvement. In 1-to-1 matching the true positive admittance rate increased for 4–6%, while the average correct 1-to-N matching improved for roughly 2–3%.

The greatest power of the method becomes apparent when the two modalities, visual and thermal, are fused. In this case the role of the glasses detection module is much more prominent, drastically decreasing the average error rate from 10% down to 3%, see Table 1. Similarly, the true admission rate increases to 74% when data is fused without special handling of glasses, and to 80% when glasses are taken into account.

5. Summary and conclusions

In this paper we described a system for personal identification based on a face biometric that uses cues from visual and thermal imagery. The two modalities are shown to complement each other, their fusion providing good illumination invariance and discriminative power between individuals. Prescription glasses, a major difficulty in the thermal spectrum, are reliably detected by our method, restricting the matching to non-affected face regions. Finally, we examined how different preprocessing methods affect recognition in the two spectra, as well as holistic and local feature-based face representations. The proposed method was shown to achieve a high recognition rate (97%) using only a small number of training images (5–7) in the presence of large illumination changes (Tables 4 and 5).

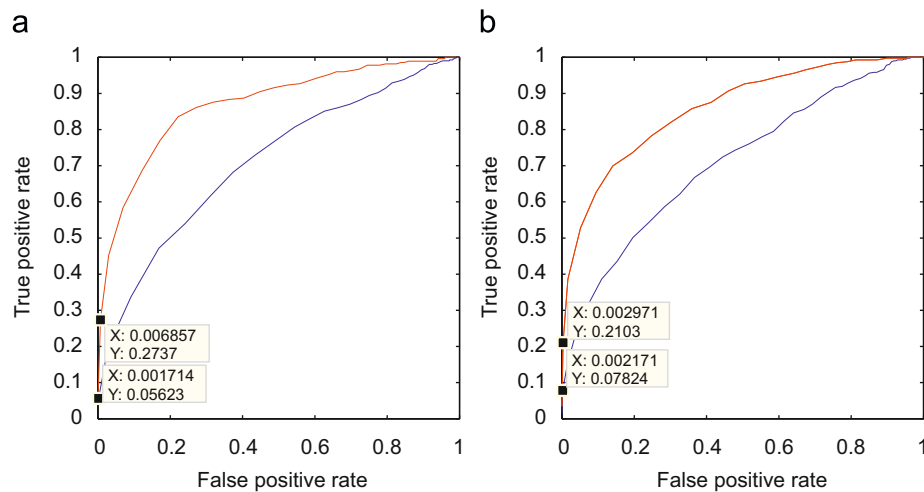


Fig. 14. Isolated local features receiver-operator characteristics (ROC): for visual (blue) and thermal (red) spectra. (a) Eyes and (b) mouth. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

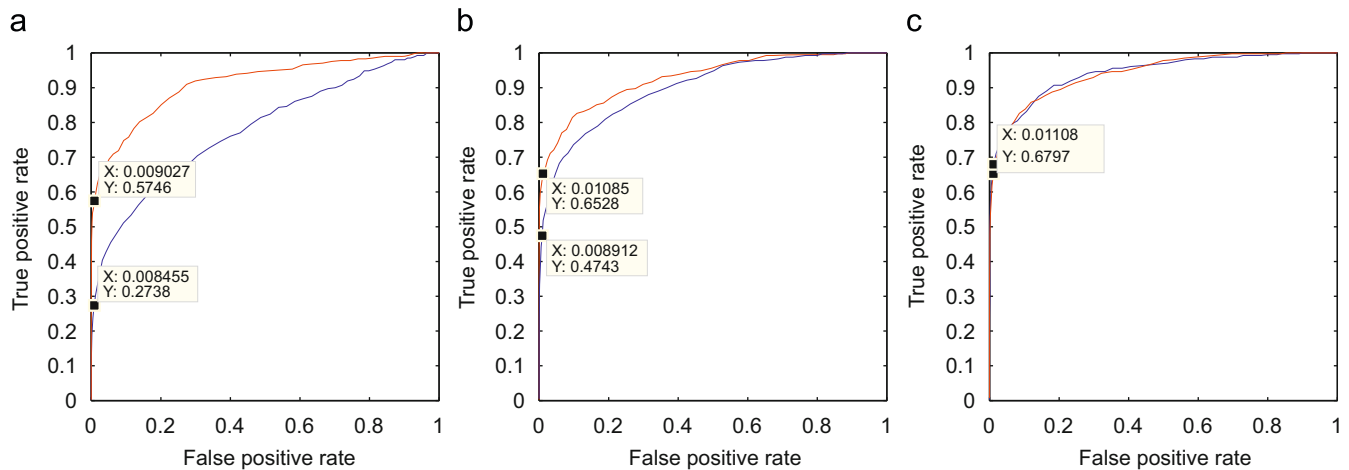


Fig. 15. Holistic representations receiver-operator characteristics (ROC) for visual (blue) and thermal (red) spectra. (a) Unprocessed, (b) band-pass filtered and (c) self-quotient image filtered. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2

A summary of the comparison of different image processing filters for 1 in 100 intruder acceptance rate.

Representation	Visual	Thermal
1% intruder acceptance		
Unprocessed/raw	0.2850	0.5803
Band-pass filtered (BP)	0.4933	0.6287
Self-quotient image (SQI)	0.6410	0.6301

Both the simple band-pass filter, and even further its locally scaled variant, greatly improve performance. This is most significant in the visual spectrum, in which image intensity in the low spatial frequency is most affected by illumination changes.

Table 3

A summary of the results for 1 in 100 intruder acceptance rate.

Representation	Visual (SQI)	Thermal (BP)
1% intruder acceptance		
Eyes	0.1016	0.2984
Mouth	0.1223	0.3037

Local features in isolation perform very poorly.

Table 4

Holistic and local features—a summary of 1-to-1 matching (verification) results.

Representation	Visual (SQI)	Thermal (BP)
1% intruder acceptance		
Holistic+eyes	0.6782	0.6499
Holistic+mouth	0.6410	0.6501
Holistic+eyes+mouth	0.6782	0.6558

Our results suggest several possible avenues for future work and improvement. A practically important issue, which was not addressed in this paper, is that of natural variation of thermal images e.g. due to post-prandial thermogenesis or increased sympathetic nervous system activity during excitement. We also intend to make further use of the thermal spectrum, by not only detecting the glasses, but also by segmenting them out. This is challenging across large pose variations, such as those contained in our test set. Another research direction we would like to pursue is that of synthetically enriching the training corpus to achieve increased robustness to pose differences between image sets (c.f.

Table 5

Feature and modality fusion—a summary of the 1-to-1 matching (verification) results.

Representation	True admission rate
1% intruder acceptance	
W/o glasses detection	0.7435
W/ glasses detection	0.8014

[68,69]). Finally, we note that empirical evaluation on a larger data set is needed to clearly identify the limitations and scalability of the proposed method.

Acknowledgments

We would like to thank Trinity College Cambridge, the Toshiba Corporation and Delphi Electronics Ltd. for their kind support for our research, the volunteers whose face videos were entered in our database and Cambridge Commonwealth Trust.

References

- [1] O. Arandjelović, R. Cipolla, Face recognition from video using the generic shape-illumination manifold, in: Proceedings of European Conference on Computer Vision (ECCV), vol. 4, 2006, pp. 27–40.
- [2] T.-J. Chin, D. Suter, A new distance criterion for face recognition using image sets, in: Proceedings of the Asian Conference on Computer Vision (ACCV), 2006, pp. 549–558.
- [3] J. Sivic, M. Everingham, A. Zisserman, Person spotting: video shot retrieval for face sets, in: Proceedings of the IEEE International Conference on Image and Video Retrieval (CIVR), 2005, pp. 226–236.
- [4] P.J. Phillips, P. Grother, R.J. Micheals, D.M. Blackburn, E. Tabassi, J.M. Bone, FRVT 2002: overview and summary, Technical Report, National Institute of Justice, March 2003.
- [5] O. Arandjelović, R. Cipolla, Incremental learning of temporally-coherent Gaussian mixture models, in: Proceedings of the IAPR British Machine Vision Conference (BMVC), vol. 2, September 2005, pp. 759–768.
- [6] T. Emidio de Campos, R.S. Feris, R.M. Cesar Jr., Eigenfaces versus eigeneyes: first steps toward performance assessment of representations for face recognition, in: Proceedings of the Mexican International Conference on Artificial Intelligence, 2000, pp. 193–201.
- [7] V. Krüger, A. Happe, G. Sommer, Affine real-time face tracking using Gabor wavelet networks, in: Proceedings of the IEEE International Conference on Pattern Recognition (ICPR), vol. 1, 2000, pp. 127–130.
- [8] V. Krüger, G. Sommer, Gabor wavelet networks for efficient head pose estimation, *Journal of the Optical Society of America* 19 (6) (2002) 1112–1119.
- [9] L. Torres, L. Lorente, J. Vila, Automatic face recognition of video sequences using self-eigenfaces, in: Proceedings of the IAPR International Symposium on Image/Video Communications over Fixed and Mobile Networks, 2000.
- [10] J.R. Price, T.F. Gee, Towards robust face recognition from video, in: Proceedings of the Applied Image Pattern Recognition Workshop, Analysis and Understanding of Time Varying Imagery, 2001, pp. 94–102.
- [11] J. Kittler, M. Hatef, R. Duin, J. Matas, On combining classifiers, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20 (3) (1998) 226–239.
- [12] G.J. Edwards, C.J. Taylor, T.F. Cootes, Learning to identify and track faces in image sequences, in: Proceedings of the IAPR British Machine Vision Conference (BMVC), 1997, pp. 130–139.
- [13] O. Yamaguchi, K. Fukui, K. Maeda, Face recognition using temporal image sequence, in: Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition (FGR), vol. (10), 1998, pp. 318–323.
- [14] R. Gittins, Canonical Analysis: A Review with Applications in Ecology, Springer, Berlin, 1985.
- [15] H. Hotelling, Relations between two sets of variates, *Biometrika* 28 (1936) 321–372.
- [16] T.F. Cootes, G.J. Edwards, C.J. Taylor, Active appearance models, in: Proceedings of the European Conference on Computer Vision (ECCV), vol. 2, 1998, pp. 484–498.
- [17] C. Wallraven, V. Blanz, T. Vetter, 3D reconstruction of faces—combining stereo with class-based knowledge, in: Proceedings of the Deutsche Arbeitsgemeinschaft für Mustererkennung (DAGM) Symposium, 1999, pp. 405–412.
- [18] L.B. Wolff, D.A. Socolinsky, C.K. Eveland, Quantitative measurement of illumination invariance for face recognition using thermal infrared imagery, in: Proceedings of the IEEE International Workshop on Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS), 2001.
- [19] D. Socolinsky, A. Selinger, J. Neuheisel, Face recognition with visible and thermal infrared imagery, *Computer Vision and Image Understanding (CVIU)* 91 (1–2) (2003) 72–114.
- [20] F. Prokoski, History, current status, and future of infrared identification, in: Proceedings of IEEE International Workshop on Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS), 2000, pp. 5–14.
- [21] D. Socolinsky, A. Selinger, Thermal face recognition in an operational scenario, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), vol. 2, 2004, pp. 1012–1019.
- [22] S. Kong, J. Heo, B. Abidi, J. Paik, M. Abidi, Recent advances in visual and infrared face recognition—a review, *Computer Vision and Image Understanding (CVIU)* 97 (1) (2005) 103–135.
- [23] D. Socolinsky, A. Selinger, Comparative study of face recognition performance with visible and thermal infrared imagery, in: Proceedings of IEEE International on Pattern Recognition (ICPR), 2002, pp. 217–222.
- [24] A. Selinger, D. Socolinsky, Appearance-based facial recognition using visible and thermal imagery: a comparative study, Technical Report 02-01, Equinox Corporation, 2002.
- [25] A. Srivastana, X. Liu, Statistical hypothesis pruning for recognizing faces from infrared images, *Image and Vision Computing* 21 (7) (2003) 651–661.
- [26] P. Buddharaju, I. Pavlidis, I. Kakadiaris, Face recognition in the thermal infrared spectrum, in: Proceedings of the IEEE International Workshop on Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS), 2004, p. 133.
- [27] G. Friedrich, Y. Yeshurun, Seeing people in the dark: face recognition in infrared images, in: Proceedings of the IAPR British Machine Vision Conference (BMVC), 2003, pp. 348–359.
- [28] C.L. Lin, K.C. Fan, Biometric verification using thermal images of palm-dorsa vein patterns, *IEEE Transactions on Circuits and Systems for Video Technology* 14 (2) (2004) 199–213.
- [29] S.K. Im, H.S. Choi, S.W. Kim, A direction-based vascular pattern extraction algorithm for hand vascular pattern verification, *ETRI Journal* 25 (2) (2003) 101–108.
- [30] T. Shimooka, K. Shimizu, Artificial immune system for personal identification with finger vein pattern, in: Proceedings of the International Conference on Knowledge-Based Intelligent Information and Engineering Systems, September 2004, pp. 511–518.
- [31] N. Miura, A. Nagasaka, T. Miyatake, Feature extraction of finger vein patterns based on iterative line tracking and its application to personal identification, *Systems and Computers in Japan* 35 (7) (2004) 61–71.
- [32] F.J. Prokoski, R. Riedel, BIOMETRICS: personal identification in networked society, in: *Infrared Identification of Faces and Body Parts*, Kluwer Academic Publishers, Dordrecht, 1998, pp. 191–211 (Chapter 9).
- [33] P. Buddharaju, I.T. Pavlidis, P. Tsiamyrtzis, Physiology-based face recognition, in: Proceedings of IEEE Conference on Advanced Video and Signal Based Surveillance (AVSS), September 2005.
- [34] O. Arandjelović, R.I. Hammoud, R. Cipolla, Multi-sensory face biometric fusion (for personal identification), in: Proceedings of IEEE International Workshop on Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS), June 2006, pp. 128–135.
- [35] S.Z. Li, R. Chu, S. Liao, L. Zhang, Illumination invariant face recognition using near-infrared images, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 29 (4) (2007) 627–639.
- [36] V. Chatzis, A.G. Bors, I. Pitas, Multimodal decision-level fusion for person authentication, *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans* 29 (6) (1999) 674–681.
- [37] S. Ben-Yacoub, Y. Abdeljaoued, E. Mayoraz, Fusion of face and speech data for person identity verification, *IEEE Transactions on Neural Networks* 10 (5) (1998) 1065–1075.
- [38] E.S. Bigun, J. Bigun, B. Duc, S. Fischer, Expert conciliation for multimodal person authentication systems using bayesian statistics, in: Proceedings of International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA), 1997, pp. 291–300.
- [39] P. Verlinde, G. Cholet, Comparing decision fusion paradigms using k-NN based classifiers, decision trees and logistic regression in a multi-modal identity verification application, in: Proceedings of International Conference on Audio, and Video-Based Biometric Person Authentication (AVBPA), vol. 5(2), 1999, pp. 188–193.
- [40] Y. Wang, T. Tan, A.K. Jain, Combining face and iris biometrics for identity verification, in: Proceedings of International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA), 2003, pp. 805–813.
- [41] R. Brunelli, D. Falavigna, Person recognition using multiple cues, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 17 (10) (1995) 955–966.
- [42] A. Ross, A.K. Jain, Information fusion in biometrics, *Pattern Recognition Letters* 24 (13) (2003) 2115–2125.
- [43] A. Ross, R. Govindarajan, Feature level fusion using hand and face biometrics, in: Proceedings of SPIE Conference on Biometric Technology for Human Identification II, vol. 5779, March 2005, pp. 196–204.
- [44] J. Heo, S. Kong, B. Abidi, M.A. Abidi, Fusion of visual and thermal signatures with eyeglass removal for robust face recognition, in: Proceedings of IEEE International Workshop on Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS), 2004, pp. 94–99.

- [45] A. Gyaourova, G. Bebis, I. Pavlidis, Fusion of infrared and visible images for face recognition, in: Proceedings of European Conference on Computer Vision (ECCV), vol. 4, May 2004, pp. 456–468.
- [46] J.G. Wang, E. Sung, R. Venkateswarlu, Registration of infrared and visible-spectrum imagery for face recognition, in: Proceedings of IEEE International Conference on Automatic Face and Gesture Recognition (FGR), May 2004, pp. 638–644.
- [47] X. Chen, P. Flynn, K. Bowyer, IR and visible light face recognition, *Computer Vision and Image Understanding (CVIU)* 99 (3) (2005) 332–358.
- [48] X. Chen, P. Flynn, K. Bowyer, Visible-light and infrared face recognition, in: Proceedings of Workshop on Multimodal User Authentication, 2003, pp. 48–55.
- [49] J. Heo, B. Abidi, S.G. Kong, M. Abidi, Performance comparison of visual and thermal signatures for face recognition, *Biometric Consortium Conference*, September 2003.
- [50] A. Ross, K. Nandakumar, A.K. Jain, *Handbook of Multibiometrics*, Springer, New York, USA, 2006.
- [51] R. Cappelli, D. Maio, D. Maltoni, A computational approach to edge detection, in: Proceedings of International Workshop on Multiple Classifier Systems, 2000, pp. 351–361.
- [52] Identix Ltd. Faceit <<http://www.Faceit.com/>>.
- [53] M. Bichsel, A.P. Pentland, Human face recognition and the face image set's topology, *Computer Vision, Graphics and Image Processing: Image Understanding* 59 (2) (1994) 254–261.
- [54] Å. Björck, G.H. Golub, Numerical methods for computing angles between linear subspaces, *Mathematics of Computation* 27 (123) (1973) 579–594.
- [55] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, second ed., Cambridge University Press, Cambridge, 1992.
- [56] O. Arandjelović, A. Zisserman, On film character retrieval in feature-length films, in: *Interactive Video: Algorithms and Technologies*, Springer, Berlin, 2006, pp. 89–103.
- [57] T.L. Berg, A.C. Berg, J. Edwards, M. Maire, R. White, Y.W. Teh, E. Learned-Miller, D.A. Forsyth, Names and faces in the news, Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), vol. 2, 2004, pp. 848–854.
- [58] D. Cristinacce, T.F. Cootes, I. Scott, A multistage approach to facial feature detection, Proceedings of IAPR British Machine Vision Conference (BMVC), vol. 1, 2004, pp. 277–286.
- [59] P.F. Felzenszwalb, D. Huttenlocher, Pictorial structures for object recognition, *International Journal of Computer Vision (IJCV)* 61 (1) (2005) 55–79.
- [60] L. Trujillo, G. Olague, R. Hammoud, B. Hernandez, E. Romero, Automatic feature localization in thermal images for facial expression recognition, *Computer Vision and Image Understanding (CVIU)* 106 (2-3) 258–269.
- [61] A. Fitzgibbon, A. Zisserman, On affine invariant clustering and automatic cast listing in movies, in: Proceedings of European Conference on Computer Vision (ECCV), 2002, pp. 304–320.
- [62] Y. Adini, Y. Moses, S. Ullman, Face recognition: the problem of compensating for changes in illumination direction, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 19 (7) (1997) 721–732.
- [63] H. Wang, S.Z. Li, Y. Wang, Face recognition under varying lighting conditions using self quotient image, in: Proceedings of IEEE International Conference on Automatic Face and Gesture Recognition (FGR), 2004, pp. 819–824.
- [64] D.S. Bolme, Elastic bunch graph matching, Master's Thesis, Colorado State University, 2003.
- [65] L. Wiskott, J.-M. Fellous, N. Krüger, C. von der Malsburg, Face recognition by elastic bunch graph matching, *Intelligent Biometric Techniques in Fingerprint and Face Recognition*, 1999, pp. 355–396.
- [66] O. Arandjelović, R. Cipolla, A methodology for rapid illumination-invariant face recognition using image processing filters, *Computer Vision and Image Understanding (CVIU)* 113 (2008).
- [67] T.C. Walker, R.K. Miller, *Health Care Business Market Research Handbook*, fifth ed., Richard K. Miller & Associates Inc, Norcross, GA, 2001.
- [68] A.M. Martinez, Recognizing imprecisely localized, partially occluded and expression variant faces from a single sample per class, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 24 (6) (2002) 748–763.
- [69] K.K. Sung, T. Poggio, Example-based learning for view-based human face detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 20 (1) (1998) 39–51.

About the Author—OGNJEN ARANDJELOVIĆ is a Research Fellow at Trinity College, Cambridge. He graduated top of his class from the Department of Engineering Science at the University of Oxford (M.Eng.). In 2007 he was awarded the Ph.D. degree from the University of Cambridge. His main research interests are computer vision and machine learning, and their application in other scientific disciplines. He is a Fellow of the Cambridge Overseas Trust and a winner of multiple best research paper awards.

About the Author—RIAD HAMMOUD is a senior research scientist at the World Headquarters of Electronics and Safety Division of Delphi Corporation. In 2001 he was awarded a PhD degree in Computer Vision and Robotics from INRIA Rhone-Alpes. After completing his doctorate, he launched a startup on the campus of Grenoble on Interactive Video Tech. His research interests include automatic target classification in and beyond the visible spectrum, eye tracking, biometrics, and real-time optimization techniques for safety, security and entertainment applications. He was nominated by US government as an outstanding researcher in May 2005.

About the Author—ROBERTO CIPOLLA received the BA degree (Engineering) from the University of Cambridge in 1984 and the MSE degree (Electrical Engineering) from the University of Pennsylvania in 1985. In 1991, he was awarded the D.Phil. degree (computer vision) from the University of Oxford. His research interests are in computer vision and robotics and include the recovery of motion and 3D shape of visible surfaces from image sequences, visual tracking and navigation, robot hand-eye coordination, algebraic and geometric invariants for object recognition and perceptual grouping, novel man-machine interfaces using visual gestures, and visual inspection. He has authored three books, edited six volumes, and coauthored more than 200 papers.