

Vesselness Features and the Inverse Compositional AAM for Robust Face Recognition using Thermal IR

Reza Shoja Ghiass[†], Ognjen Arandjelović[‡], Hakim Bendada[†] and Xavier Maldague[†]

[†] Université Laval, Quebec, Canada

[‡] Deakin University, Geelong, Australia

Abstract

Over the course of the last decade, infrared (IR) and particularly thermal IR imaging based face recognition has emerged as a promising complement to conventional, visible spectrum based approaches which continue to struggle when applied in the real world. While inherently insensitive to visible spectrum illumination changes, IR images introduce specific challenges of their own, most notably sensitivity to factors which affect facial heat emission patterns, e.g. emotional state, ambient temperature, and alcohol intake. In addition, facial expression and pose changes are more difficult to correct in IR images because they are less rich in high frequency detail which is an important cue for fitting any deformable model. In this paper we describe a novel method which addresses these major challenges. Specifically, to normalize for pose and facial expression changes we generate a synthetic frontal image of a face in a canonical, neutral facial expression from an image of the face in an arbitrary pose and facial expression. This is achieved by piecewise affine warping which follows active appearance model (AAM) fitting. This is the first publication which explores the use of an AAM on thermal IR images; we propose a pre-processing step which enhances detail in thermal images, making AAM convergence faster and more accurate. To overcome the problem of thermal IR image sensitivity to the exact pattern of facial temperature emissions we describe a representation based on reliable anatomical features. In contrast to previous approaches, our representation is not binary; rather, our method accounts for the reliability of the extracted features. This makes the proposed representation much more robust both to pose and scale changes. The effectiveness of the proposed approach is demonstrated on the largest public database of thermal IR images of faces on which it achieved 100% identification rate, significantly outperforming previously described methods.

Introduction

Notwithstanding the substantial and continued research effort, the practical success of face recognition using conventional imaging equipment has been limited. In an effort to overcome some of the key challenges which remain, such as appearance changes due to varying illumination and disguises (including facial wear and hair), the use of imaging outside the visible part of the electromagnetic spectrum

(wavelength approximately in the range 390 – 750 nm) has been explored. Most notably, the use of infrared data (wavelength approximately in the range 0.74 – 300 μm) has been increasingly popular, in no small part due to the reducing cost of infrared cameras.

In the literature, it has been customary to divide the IR spectrum into four sub-bands: near IR (NIR; wavelength 0.75 – 1.4 μm), short wave IR (SWIR; wavelength 1.4 – 3 μm), medium wave IR (MWIR; wavelength 3 – 8 μm), and long wave IR (LWIR; wavelength 8 – 15 μm). This division of the IR spectrum is also observed in the manufacturing of IR cameras, which are often made with sensors that respond to electromagnetic radiation constrained to a particular sub-band. It should be emphasized that the division of the IR spectrum is not arbitrary. Rather, different sub-bands correspond to continuous frequency chunks of the solar spectrum which are divided by absorption lines of different atmospheric gasses (Maldague 2001). In the context of face recognition, one of the largest differences between different IR sub-bands emerges as a consequence of the human body's heat emission spectrum. Specifically, most of the heat energy is emitted in LWIR sub-band, which is why it is often referred to as the thermal sub-band (this term is often extended to include the MWIR sub-band). Significant heat is also emitted in the MWIR sub-band. Both of these sub-bands can be used to *passively* sense facial thermal emissions without an external source of light. This is one of the reasons why LWIR and MWIR sub-bands have received the most attention in the face recognition literature. In contrast to them, facial heat emission in the SWIR and NIR sub-bands is very small and systems operating on data acquired in these sub-bands require appropriate illuminators i.e. recognition is *active* in nature. In recent years, the use of NIR also started received increasing attention from the face recognition community, while the utility of the SWIR sub-band has yet to be studied in depth. An up to date, thorough literature review can be found in (Ghiass et al. 2013b).

Advantages of IR The foremost advantage of IR data in comparison with conventional visible spectrum images in the context of face recognition lies in its invariance to visible spectrum illumination. This is inherent in the very nature of IR imaging and, considering the challenge that variable illumination conditions present to face recogni-

tion systems, a major advantage. What is more, IR energy is also less affected by scattering and absorption by smoke or dust than reflected visible light (Chang et al. 2008; Nicolo and Schmid 2011). Unlike visible spectrum imaging, IR imaging can be used to extract not only exterior but also useful subcutaneous anatomical information, such as the vascular network¹ of a face (Buddharaju et al. 2007) or its blood perfusion patterns (Wu et al. 2007). Finally, thermal vision can be used to detect facial disguises (Pavlidis and Symosek 2000) as well.

Challenges in IR The use of IR images for AFR is not void of its problems and challenges. For example, MWIR and LWIR images are sensitive to the environmental temperature, as well as the emotional, physical and health condition of the subject. They are also affected by alcohol intake. Another potential problem is that eyeglasses are opaque to the greater part of the IR spectrum (LWIR, MWIR and SWIR) (Tasman and Jaeger 2009). This means that a large portion of the face wearing eyeglasses may be occluded, causing the loss of important discriminative information. The complementary nature of visible spectrum images in this regard has inspired various multi-modal fusion methods (Heo et al. 2004; Arandjelović, Hammoud, and Cipolla 2006). Another consideration of interest pertains to the impact of sunlight if recognition is performed outdoors and during daytime. Although invariant to the changes in the illumination by visible light itself (by definition), the IR “appearance” in the NIR and MWIR sub-bands is affected by sunlight which has significant spectral components at the corresponding wavelengths. This is one of the key reasons why NIR and SWIR based systems which perform well indoors struggle when applied outdoors (Li et al. 2007a).

Previous work The earliest attempts at examining the potential of infrared imaging for face recognition date back to the work done in (Prokoski, Riedel, and Coffin 1992). Most of the automatic methods which followed closely mirrored the methods developed for visible spectrum based recognition. Generally, these used holistic face appearance in a simple statistical manner, with little attempt to achieve any generalization, relying instead on the availability of training data with sufficient variability of possible appearance for each subject (Cutler 1996; Socolinsky et al. 2001; Selinger and Socolinsky 2004). More sophisticated holistic approaches recently investigated include statistical models based on Gaussian mixtures (Elguebaly and Bouguila 2011) and compressive sensing (Lin et al. 2011). Numerous feature based approaches have also been described. The use of locally binary patterns was proposed in (Li et al. 2007b)

¹It is important to emphasize that none of the existing publications on face recognition using ‘vascular network’ based representations provide any evidence that the extracted structures are indeed blood vessels. Thus the reader should understand that we use this term for the sake of consistency with previous work, and that we do *not* claim that what we extract in this paper is an actual vascular network. Rather we prefer to think of our representation as a *function* of the underlying vasculature.

and (Goswami et al. 2011), SIFT features in (Maeng et al. 2011), wavelets in (Srivastana and Liu 2003) and (Nicolo and Schmid 2011), and curvelets in (Xie et al. 2009). The method in (Wu et al. 2005) is one of the few in the literature which attempts to extract useful subcutaneous information from IR appearance. Using a blood perfusion model Wu *et al.* infer the blood perfusion pattern corresponding to an IR image of a face. In (Buddharaju, Pavlidis, and Tsiamyrtzis 2005) the vascular network of a face is extracted instead. Buddharaju *et al.* represent vascular networks as a binary images (each pixel either is or is not a part of the network) and match them using a method adopted from fingerprint recognition: using salient loci of the networks (such as bifurcation points). While successful in the recognition of fingerprints which are planar, this approach is not readily adapted to deal with pose changes expected in many practical applications of face recognition. In addition, as we will explain in further detail in ‘Multi-scale blood vessel extraction’, the binary nature of their method for vascular network extraction makes it very sensitive to face scale and image resolution.

Method details

Our algorithm comprises two key components. The first of these involves the extraction of an image based face representation, which is largely invariant to changes in facial temperature. This invariance is achieved by focusing on the extraction of anatomical information of a face, rather than absolute or relative temperature. The other component of our algorithm normalizes the changes in the person’s pose by creating a synthetic IR image of the person facing the camera in a canonical, neutral facial expression. A detailed explanation of the two methods, as well as different preprocessing steps necessary to make the entire system automatic, is explained in detail next.

Pose normalization using the AAM

Much like in the visible spectrum, the appearance of face in the IR spectrum is greatly affected by the person’s head pose relative to the camera (Friedrich and Yeshurun 2003). Therefore it is essential that recognition is performed either using features which are invariant to pose changes or that pose is synthetically normalized. In the present paper we adopt the latter approach. Specifically, from an input image of a face in an initially unknown pose (up to 30° yaw difference from frontal) we synthetically generate a frontal, camera facing image of the person using an active appearance model (AAM) (Cootes, Edwards, and Taylor 1998). Although widely used for visible spectrum based face recognition (Feng et al. 2011; Sauer, Cootes, and Taylor 2011), to the best of the knowledge of these authors, this is the first published work which has attempted to apply it on IR data (specifically thermal IR data).

Face segmentation One of the key problems encountered in practical application of AAM is their initialization. If the initial geometric configuration of model is too far from the correct solution, the fitting may converge to a locally optimal but incorrect set of values (loci of salient points).

This is a particularly serious concern for thermal IR images since thermal appearance of faces is less rich in detail which guides and constrains the AAM. For this reason, in our method face segmentation is performed first. This accomplishes two goals. Firstly, the removal of any confounding background information helps the convergence in the fitting of the model. Secondly, the AAM can be initialized well, by considering the shape and scale of the segmented foreground region.

Unlike in the visible spectrum, in which background clutter is often significant and in which face segmentation can be a difficult task, face segmentation in thermal IR images is in most cases far simpler. Indeed, in the present paper we accomplish the bulk of work using simple thresholding. We create a provisional segmentation map by declaring all pixels with values within a range between two thresholds, T_{low} and T_{up} , as belonging to the face region (i.e. foreground) and all others as background. An example is shown in Fig. 1(a). The provisional map is further refined by performing morphological opening and closing operations, using a circular structuring element (we used a circle whose approximate area is 6% of the area of the segmented face ellipse). This accomplishes the removal of spurious artefacts which may occur at the interface between the subject’s skin and clothing for example, and correctly classifies facial skin areas which are unusually cold or hot (e.g. the exposure to cold surroundings can transiently create cold skin patches in the regions of poor perifocal vascularity). An example of the final segmentation mask is shown in Fig. 1(c) and the segmented image output in Fig. 1(d).



Figure 1: The first step in the proposed algorithm is to segment out the face. This removed image areas unrelated to the subject’s identity and aids in the convergence of the AAM. From (a) the original image (b) provisional segmentation mask is created using temperature thresholding, after which (c) morphological operators are used to increase the segmentation accuracy to outliers e.g. such as which occur at the interface of facial and non-facial regions, producing (d) the final result with the background correctly suppressed.

Detail enhancement As mentioned earlier, thermal IR images of faces are much less rich in fine detail than visible spectrum images (Chen, Flynn, and Bowyer 2003), which be readily observed in the example in Fig. 2(a). This makes the problem of AAM fitting all the more challenging. For this reason, we do not train or fit the AAM on segmented thermal images themselves, but on processed and detail enhanced images. In our experiments we found that the additional detail the proposed filtering brings out greatly aids in guiding the AAM towards the correct solution.

The method for detail enhancement we propose is a form of non-linear high pass filtering. Firstly, we anisotropically

diffuse the input image. We adopt the form of the diffusion:

$$\frac{\partial I}{\partial t} = \nabla \cdot (c(\|\nabla I\|) \nabla I) = \nabla c \cdot \nabla I + c(\|\nabla I\|) \Delta I, \quad (1)$$

with the diffusion parameter c constant over time (i.e. filtering iterations) but spatially varying and dependent on the magnitude of the image gradient:

$$c(\|\nabla I\|) = \exp \left\{ -\frac{\|\nabla I\|}{k^2} \right\}. \quad (2)$$

We used $k = 20$. The detail enhanced image is then computed by subtracting the diffused image from the original: $I_e = I - I_d$. This is illustrated on a typical example in Fig. 2(a,b).




Figure 2: The AAM is notoriously sensitive to initialization. This potential problem is even greater when the model is used on thermal images, which lack characteristic, high frequency content. We increase the accuracy of AMM fitting first by (a) creating an anisotropically smoothed thermal image, which is then (b) subtracted from the original image to produce an image with enhanced detail.

Inverse compositional AAM fitting The type of an active appearance model we are interested in here separately models the face shape, as a piecewise triangular mesh, and face appearance, covered by the mesh (Cootes, Edwards, and Taylor 1998). The model is trained using a data corpus of faces which has salient points manually annotated and which become the vertices of the corresponding triangular mesh. These training meshes are used to learn the scope of variation of shape (i.e. locations of vertices) and appearance of individual mesh faces, both using principal component analysis i.e. by retaining the first principal components as the basis of the learnt generative model. The model is applied on a novel face by finding a set of parameters (shape and appearance principal component weights) such that the difference between the corresponding piecewise affine warped image and the model predicted appearance is minimized. Formally, the fitting error can be written as:

$$e_{aam} = \sum_{\text{All pixels } \mathbf{x}} \left[A_0(\mathbf{x}) + \sum_{i=1}^m \alpha_i A_i(\mathbf{x}) - I_e(\mathbf{W}(\mathbf{x}; \mathbf{p})) \right]^2 \quad (3)$$

where \mathbf{x} are pixel loci, A_i ($i = 0 \dots m$) the retained appearance principal components and $\mathbf{W}(\mathbf{x}; \mathbf{p})$ the location of the pixel warped using the shape parameters \mathbf{p} . For additional detail, the reader is referred to the original publication (Cootes, Edwards, and Taylor 1998).

The most straightforward approach to minimizing the error described in Eq. 3 is by using gradient descent. However,

this is slow. A popular alternative proposed in (Cootes, Edwards, and Taylor 1998) uses an approximation that there is a linear relationship between the fitting error term in Eq. 3, and the updates to the shape and appearance parameters, respectively $\Delta \mathbf{p}$ and $\Delta \alpha_i$ (more complex variations on the same theme include (Sclaroff and Isidoro 1998; Cootes, Edwards, and Taylor 2001)). What is more, the linear relationship is assumed not to depend on the model parameters, facilitating a simple learning of the relationship from the training data corpus. While much faster than gradient descent, this fitting procedure has been shown to produce inferior fitting results in comparison to the approach we adopt here, the inverse compositional AAM (ICAAM) (Matthews and Baker 2004). Our experiments show that the advantages of ICAAM are even greater when fitting is done on thermal IR images, as opposed to visual ones (used by all of the aforementioned authors).

There are two key differences between the conventional AAM fitting and ICAAM. Firstly, instead of estimating a simple update of the parameters $\Delta \mathbf{p}$, the compositional AAM algorithm estimates the update to the warp itself, i.e. $\mathbf{W}(\mathbf{x}, \mathbf{p})$. This particular idea was first advanced in (Lucas and Kanade 1981). Secondly, in the inverse compositional algorithm, the direction of piecewise linear warping is inverted. Instead of warping the input image to fit the reference mesh, the warping is performed in the opposite direction. In other words, the error minimized becomes:

$$e_{icaam} = \sum_{\text{All pixels } \mathbf{x}} [I_e(\mathbf{W}(\mathbf{x}; \mathbf{p})) - A_0(\mathbf{W}(\mathbf{x}; \mathbf{p}))]^2. \quad (4)$$

While it can be shown that the minimization of Eq. 4 is equivalent to that of Eq. 3, this inverse formulation allows much of the intensive computation to be pre-computed, thus resulting in faster convergence in addition to the increased fitting accuracy provided by the compositional AAM. A typical result of applying our approach is shown in Fig. 3.




Figure 3: (a) A converged AAM mesh is superimposed on the segmented and enhanced thermal image, and (b) a synthetic image of a frontal face generated by piecewise affine warping of the original image.

Multi-scale blood vessel extraction

Following the application of piecewise affine warping of the input image using ICAAM, all training and query faces are normalized for pose variations and the corresponding synthetically generated images contain frontal faces. Our goal in this stage of our algorithm is to extract from these images a person-specific representation which is firstly invariant to

temperature changes and secondly robust to small imperfections of the preceding pose normalization.

Differences from previous work As already argued in ‘Introduction’, the absolute temperature value of a particular point on a person’s face can greatly vary as the conditions in which IR data is acquired are changed. The *relative* temperature of different regions across the face is equally variable – even simple physiological changes, such as an increase in sympathetic nervous system activity are effected non-uniformly. These observations strongly motivate the development of representations which are based on invariable anatomical features, which are unaffected by the aforementioned changeable distributions in heat emission.

The use of a vascular anatomical invariant was proposed in (Buddharaju et al. 2007). The key underlying observation is that blood vessels are somewhat warmer than the surrounding tissues, allowing them to be identified in thermograms. These temperature differences are very small and virtually imperceptible to the naked eye, but inherently maintained regardless of the physiological state of the subject. An important property of vascular networks which makes them particularly attractive for use in recognition is that the blood vessels are ‘hardwired’ at birth and form a pattern which remains virtually unaffected by factors such as aging, except for predictable growth (Persson and Buschmann 2011).

In the present work we too adopt a vascular network based approach, but with several important differences in comparison with the previously proposed methods. The first of these is to be found in the manner vascular structures are extracted. Buddharaju *et al.* adopt the use of a simple image processing filter based on the ‘top hat’ function. We found that the output produced by this approach is very sensitive to the scale of the face (something not investigated by the original authors) and thus lacks robustness to the distance of the user from the camera. In contrast, the approach taken in the present paper is specifically aimed at extracting vessel-like structures, and it does so in a multi-scale fashion, integrating evidence across scales, see Fig. 4. The second major difference lies in the form in which the extracted network is represented. Buddharaju *et al.* produce a binary image, in which each pixel is deemed either as belonging to the network or not, without any accounting for the uncertainty associated with this classification. This aspect of their approach makes it additionally sensitive to variable pose and temperature changes across the face; see (Chen et al. 2011) for related criticism. In contrast, in our baseline representation each pixel is real-valued, its value quantifying the degree of confidence ($\in [0, 1]$) that it is a part of the vascular structure.

Vesselness We extract characteristic anatomical features of a face from its thermal image using the method proposed in (Frangi et al. 1998). Their so-called vesselness filter, first proposed for the use on 3D MRI data, extracts tubular structures from an image. For a 2D image consider the two eigenvalues λ_1 and λ_2 of the Hessian matrix computed at a certain image locus and at a particular scale. Without loss of generality let us also assume that $|\lambda_1| \leq |\lambda_2|$. The two key values used to quantify how tubular the local structure at this scale is are $\mathcal{R}_A = |\lambda_1|/|\lambda_2|$ and $\mathcal{S} = \sqrt{\lambda_1^2 + \lambda_2^2}$. The former of these measures the degree of local ‘blobiness’.

Vascular network of Buddhharaju *et al.*

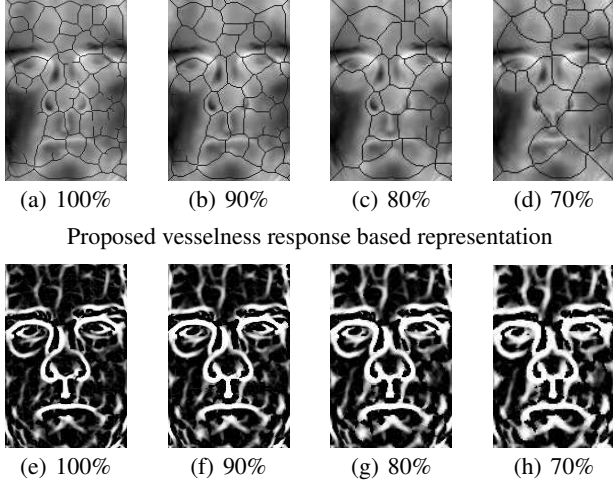


Figure 4: One of the major limitations of the approach proposed by Buddhharaju *et al.* lies in its ‘crisp’ binary nature: a particular pixel is deemed either a part of the vascular network or not. A consequence is that the extracted vascular network is highly sensitive to the scale of the input image (and thus to the distance of the user from the camera as well as the spatial resolution of the camera). (a-d) Even small changes in face scale can effect large topological changes on the result (note that the representation of the interest is the vascular network, shown in black, which is only superimposed on the images it is extracted from for the benefit of the reader). (e-h) In contrast, the proposed vesselness response based representation encodes the certainty that a particular pixel locus is a reliable vessel pattern, and exhibits greater resilience to scale changes.

If the local appearance is blob-like, the Hessian is approximately isotropic and $|\lambda_1| \approx |\lambda_2|$ making \mathcal{R}_A close to 1. For a tubular structure \mathcal{R}_A should be small. On the other hand, \mathcal{S} ensures that there is sufficient local information content at all: in nearly uniform regions, both eigenvalues of the corresponding Hessian will have small values. For a particular scale of image analysis s , the two measures, \mathcal{R}_A and \mathcal{S} , are then unified into a single vesselness measure:

$$\mathcal{V}(s) = \begin{cases} 0 & \text{if } \lambda_2 > 0 \\ e^{-\frac{\mathcal{R}_A}{2\alpha^2}} \times (1 - e^{-\frac{\mathcal{S}}{2c^2}}) & \text{otherwise,} \end{cases} \quad (5)$$

where α and c are the parameters that control the sensitivity of the filter to \mathcal{R}_A and \mathcal{S} . Finally, if an image is analyzed across scales from s_{min} to s_{max} , the vesselness of a particular image locus can be computed as the maximal vesselness across the range:

$$\mathcal{V}_0 = \max_{s_{min} \leq s \leq s_{max}} \mathcal{V}(s) \quad (6)$$

Vesselness at three different scales for an example thermal image is illustrated in Fig. 5(a-c), and the corresponding multi-scale result in Fig. 5(d).

Matching After the fitting of an AAM to a gallery or novel face, all face images, or indeed the extracted vesselness signatures, can be warped to a canonical frame. This warping normalizes data with respect to pose and facial expression, while the underlying representation ensures robustness with

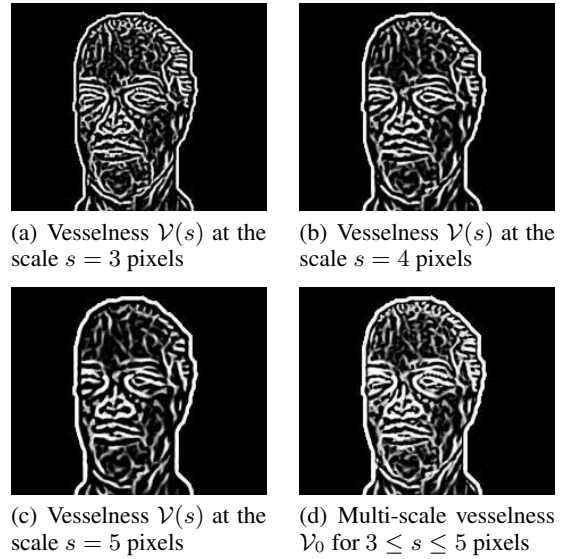


Figure 5: (a-c) The output of the vesselness filter at three different scales, and (d) the corresponding integrated multi-scale result.

respect to various extrinsic factors which affect the face temperature (absolute as well as relative across different parts of the face, as discussed in ‘Multi-scale blood vessel extraction’). Since the vesselness image inherently exhibits small shift invariance due to the smoothness of the vesselness filter response, as readily observed on an example in Fig. 5(c), thus normalized images can be directly compared. In this paper we adopt the simple cross-correlation coefficient as a measure of similarity. If I_{n1} and I_{n2} are two normalized images (warped vesselness signatures), their similarity ρ is computed as:

$$\rho(I_{n1}, I_{n2}) = \frac{\sum_{i,j} (I_{n1}(i,j) - \bar{I}_{n1})(I_{n2}(i,j) - \bar{I}_{n2})}{\sqrt{\sum_{i,j} (I_{n1}(i,j) - \bar{I}_{n1})^2 \times \sum_{i,j} (I_{n2}(i,j) - \bar{I}_{n2})^2}} \quad (7)$$

where \bar{I}_{n1} and \bar{I}_{n2} are the mean values of the corresponding images.

Evaluation

In this section we report our empirical evaluation of the methods proposed in this paper. We start by describing the data set used in our experiments, follow by an explanation of the adopted evaluation protocol, and finish with a report of the results and their discussion.

University of Houston data set We chose to use the University of Houston data set for our experiments. There are several reasons for this choice. Firstly, this is one of the largest data sets of thermal images of faces; it contains data from a greater number of individuals than Equinox (Heo *et al.* 2004), IRIS (Arandjelović, Hammoud, and Cipolla 2010) or Florida State University (Srivastana and Liu 2003) collections, and a greater variability in pose and expression than those of University of Notre Dame (Chen, Flynn, and

Bowyer 2005) or the University of California/Irvine (Pan et al. 2005). Secondly, we wanted to make our results directly comparable to those of Buddharaju *et al.* whose method bears the most resemblance to ours in spirit (but not in technical detail).

The University of Houston data set consists of a total of 7590 thermal images of 138 subjects, with a uniform number of 55 images per subject. The ethnicity, age and sex of subjects vary across the database. With the exception of four subjects, from whom data was collected in two sessions six months apart, the data for a particular subject was acquired in a single session. The exact protocol which was used to introduce pose and expression variability in the data set was not described by the authors (Buddharaju et al. 2007). Example images are shown in Fig. 6. The database is available free of charge upon request.




Figure 6: False colour thermal appearance images of a subject in the five key poses in the University of Houston data set.

AAM training We trained the AAM using 90 manually annotated images. Specifically, we annotated 30 images with different facial expression for 3 different poses: approximately frontal, and at approximately $\pm 30^\circ$ and $\pm 15^\circ$ yaw change from frontal. Additionally, we exploited the vertical symmetry of faces by including examples synthetically generated by mirroring the manually annotated images, giving the total of 180 images. Following the application of PCA on the corresponding shapes and appearances, we retain the dominant components which explain 99% of the variation. This results in a model with 85 shape and 46 appearance components.

Evaluation methodology We evaluated the proposed algorithm in a setting in which the algorithm is trained using only a single image in an arbitrary pose and facial expression. The querying of the algorithm using a novel face is also performed using a single image, in a different pose and/or facial expression. Both pose and facial expression changes present a major challenge to the current state of the art, and the consideration of the two in combination make our evaluation protocol extremely challenging (indeed, more so than any attempted by previous work), and, importantly, representative of the conditions which are of interest in a wide variety of practical applications.

Results and discussion

To assess the performance of the proposed method, we first examined its rank- N (for $N \geq 1$) recognition rate i.e. its cumulative match characteristics. Our method was found to exhibit perfect performance at rank-1 already, correctly recognizing all of the subjects in the database. This result

is far superior to the previously proposed thermal minutia points based approach (Buddharaju, Pavlidis, and Tsiamirtzis 2006) which correctly recognizes 82.5% of the individuals at rank-1 and does not reach 100% even for rank-20 matching, or indeed the iterative vascular network registration based method (Pavlidis and Buddharaju 2009) which correctly recognizes 96.2% of the individuals rank-1 and which also fails to achieve 100% even at rank-20.

The performance of the proposed method is all the more impressive when it is considered that unlike the aforementioned previous work, we perform recognition using a single training image only, and across pose and facial expression changes. Both Buddharaju *et al.*, and Pavlidis and Buddharaju train their algorithm using multiple images. In addition, it should be emphasized that they do not consider varying pose – the pose of an input face is first categorized by pose and then matched with training faces in that pose only. In contrast, we perform training using a single image only, truly recognizing *across* pose variations. A comparative summary is shown in Tab. 1.

Summary and future work

In this paper we described a novel method for face recognition using thermal IR images. Our work addressed two main challenges. These are the variations of thermal IR appearance due to (i) change in head pose and expression, and (ii) facial heat pattern emissions (e.g. affected by ambient temperature or sympathetic nervous system activity). We normalize pose and facial expression by generating a synthetic frontal image of a face, following the fitting of an AAM. Our work is the first to consider the use of AAMs on thermal IR data; we show how AAM convergence problems associated with the lack of high frequency detail in thermal images can be overcome by a pre-processing stage which enhances such detail. To achieve robustness to changes in facial heat pattern emissions, we propose a representation which is not based on either absolute or relative facial temperature but instead unchangeable anatomic features in the form of a subcutaneous vascular network. We describe a more robust vascular network extraction than that used in the literature to date. Our approach is based on the so-called vesselness filter. This method allows us to process the face in multi-scale fashion and account for the confidence that a particular image locus correspond to a vessel, thus achieving greater resilience to head pose changes, face scale and input image resolution. The effectiveness of the proposed algorithm was demonstrated on the largest publicly available data set, which includes large pose and facial expression variation. In the immediate future, our future work will concentrate on extending the method to deal with the full range of head pose variation (from frontal to profile), e.g. using an ensemble of pose-specific AAMs; see (Ghiass et al. 2013a).

References

- Arandjelović, O.; Hammoud, R. I.; and Cipolla, R. 2006. Multi-sensory face biometric fusion (for personal identification). *In Proc. IEEE International Workshop on Object Tracking and Classification Beyond the Visible Spectrum* 128–135.

Table 1: A summary of the key evaluation results and method features of the proposed algorithm, and the two previously proposed vascular network based approaches (Buddharaju, Pavlidis, and Tsiamyrtzis 2006) and (Pavlidis and Buddharaju 2009). Legend: ● large degree of invariance; ◐ some degree of invariance; ○ little to no invariance.

	Recognition rate			Invariance to			
	Rank-1	Rank-3	Rank-5	Expression	Pose	Physiological condition	Scale
AAM + multi-scale vesselness (the proposed method)	100.0%	100.0%	100.0%	●	●	●	●
vascular network alignment (Pavlidis and Buddharaju 2009)	96.2%	98.3%	99.0%	◐	◐	●	○
thermal minutiae points (Buddharaju, Pavlidis, and Tsiamyrtzis 2006)	82.5%	92.2%	94.4%	◐	◐	●	○

Arandjelović, O.; Hammoud, R. I.; and Cipolla, R. 2010. Thermal and reflectance based personal identification methodology in challenging variable illuminations. *Pattern Recognition* 43(5):1801–1813.

Buddharaju, P.; Pavlidis, I. T.; Tsiamyrtzis, P.; and Bazakos, M. 2007. Physiology-based face recognition in the thermal infrared spectrum. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(4):613–626.

Buddharaju, P.; Pavlidis, I. T.; and Tsiamyrtzis, P. 2005. Physiology-based face recognition. In *Proc. IEEE Conference on Advanced Video and Singal Based Surveillance* 354–359.

Buddharaju, P.; Pavlidis, I.; and Tsiamyrtzis, P. 2006. Pose-invariant physiological face recognition in the thermal infrared spectrum. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops* 53–60.

Chang, H.; Koschan, A.; Abidi, M.; Kong, S. G.; and Won, C. 2008. Multispectral visible and infrared imaging for face recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops* 1–6.

Chen, L.; Ju, Y.; Ding, S.; and Liu, X. 2011. Topological vascular tree segmentation for retinal images using shortest path connection. In *Proc. IEEE International Conference on Image Processing* 2137–2140.

Chen, X.; Flynn, P.; and Bowyer, K. 2003. Visible-light and infrared face recognition. In *Proc. Workshop on Multimodal User Authentication* 48–55.

Chen, X.; Flynn, P.; and Bowyer, K. 2005. IR and visible light face recognition. *Computer Vision and Image Understanding* 99(3):332–358.

Cootes, T. F.; Edwards, G. J.; and Taylor, C. J. 1998. Active appearance models. In *Proc. European Conference on Computer Vision* 2:484–498.

Cootes, T. F.; Edwards, G. J.; and Taylor, C. J. 2001. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(6):681–685.

Cutler, R. 1996. Face recognition using infrared images and eigenfaces. *Technical Report, University of Maryland*.

Elguebaly, T., and Bouguila, N. 2011. A Bayesian method for infrared face recognition. *Machine Vision Beyond Visible Spectrum* 1:123–138.

Feng, X.; Shen, X.; Zhou, M.; Zhang, H.; and Kim, J. 2011. Robust facial expression tracking based on composite constraints AAM. In *Proc. IEEE International Conference on Image Processing* 3045–3048.

Frangi, A. F.; Niessen, W. J.; Vincken, K. L.; and Viergever, M. A. 1998. Multiscale vessel enhancement filtering. *Medical Image Computing and Computer-Assisted Intervention* 130–137.

Friedrich, G., and Yeshurun, Y. 2003. Seeing people in the dark: Face recognition in infrared images. In *Proc. British Machine Vision Conference* 348–359.

Ghiass, R. S.; Arandjelović, O.; Bendada, A.; and Maldague, X. 2013a. Illumination-invariant face recognition from a single image across extreme pose using a dual dimension AAM ensemble in the thermal infrared spectrum. *International Joint Conference on Neural Networks*.

Ghiass, R. S.; Arandjelović, O.; Bendada, A.; and Maldague, X. 2013b. Infrared face recognition: a literature review. *International Joint Conference on Neural Networks*.

Goswami, D.; Chan, C. H.; Windridge, D.; and Kittler, J. 2011. Evaluation of face recognition system in heterogeneous environments (visible vs NIR). In *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops* 2160–2167.

Heo, J.; Kong, S. G.; Abidi, B. R.; and Abidi, M. A. 2004. Fusion of visual and thermal signatures with eyeglass removal for robust face recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops* 122.

Li, S.; Chu, R.; Liao, S.; and Zhang, L. 2007a. Illumination invariant face recognition using near-infrared images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(4):627–639.

Li, S. Z.; Chu, R.; Liao, S.; and Zhang, L. 2007b. Illumination invariant face recognition using near-infrared images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(4):627–639.

Lin, Z.; Wenrui, Z.; Li, S.; and Zhijun, F. 2011. Infrared face recognition based on compressive sensing and PCA. In *Proc. International Conference on Computer Science and Automation Engineering* 2:51–54.

Lucas, B., and Kanade, T. 1981. An iterative image registration technique with an application to stereo vision. In *Proc. International Joint Conference on Artificial Intelligence* 1:674–679.

Maeng, H.; Choi, H.-C.; Park, U.; Lee, S.-W.; and Jain, A. K. 2011. NFRAD: Near-infrared face recognition at a distance. In *Proc. International Joint Conference on Biometrics* 1–7.

Maldague, X. 2001. *Theory and Practice of Infrared Technology for Non Destructive Testing*. John-Wiley & Sons.

Matthews, I., and Baker, S. 2004. Active appearance models revisited. *International Journal of Computer Vision* 60(2):135–164.

Nicolo, F., and Schmid, N. A. 2011. A method for robust multi-

spectral face recognition. *In Proc. IEEE International Conference on Image Analysis and Recognition* 2:180–190.

Pan, Z.; Healey, G.; Prasad, M.; and Tromberg, B. 2005. Multi-band and spectral eigenfaces for face recognition in hyperspectral images. *In Proc. SPIE* 5779:144–151.

Pavlidis, I., and Buddharaju, P. 2009. The imaging issue in an automatic face/disguise detection system. *In Proc. IEEE Conference on Computer Vision and Pattern Recognition* 128–135.

Pavlidis, I., and Symosek, P. 2000. The imaging issue in an automatic face/disguise detection system. *In Proc. IEEE Workshop on Computer Vision Beyond the Visible Spectrum* 15–24.

Persson, A. B., and Buschmann, I. R. 2011. Vascular growth in health and disease. *Frontiers of Molecular Neuroscience* 4(14).

Prokoski, F. J.; Riedel, R. B.; and Coffin, J. S. 1992. Identification of individuals by means of facial thermography. *In Proc. IEEE International Carnahan Conference on Security Technology (ICCST): Crime Countermeasures* 120–125.

Sauer, P.; Cootes, T.; and Taylor, C. 2011. Accurate regression procedures for active appearance models. *In Proc. British Machine Vision Conference*.

Sciaroff, S., and Isidoro, J. 1998. Active blobs. *In Proc. IEEE International Conference on Computer Vision* 1146–1153.

Selinger, A., and Socolinsky, D. 2004. Face recognition in the dark. *In Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops* 8:129.

Socolinsky, D. A.; Wolff, L. B.; Neuheisel, J. D.; and Eveland, C. K. 2001. Illumination invariant face recognition using thermal infrared imagery. *In Proc. IEEE Conference on Computer Vision and Pattern Recognition* 1:527.

Srivastana, A., and Liu, X. 2003. Statistical hypothesis pruning for recognizing faces from infrared images. *Image and Vision Computing* 21(7):651–661.

Tasman, W., and Jaeger, E. A. 2009. *Duane's Ophthalmology*. Lippincott Williams & Wilkins.

Wu, S.; Song, W.; Jiang, L.; Xie, S.; Pan, F.; Yau, W.; and Ranganath, S. 2005. Infrared face recognition by using blood perfusion data. *In Proc. International Conference on Audio- and Video-Based Biometric Person Authentication* 320–328.

Wu, S. Q.; Wei, L. Z.; Fang, Z. J.; Li, R. W.; and Ye, X. Q. 2007. Infrared face recognition based on blood perfusion and sub-block DCT in wavelet domain. *In Proc. International Conference on Wavelet Analysis and Pattern Recognition* 3:1252–1256.

Xie, Z.; Wu, S.; Liu, G.; and Fang, Z. 2009. Infrared face recognition based on radiant energy and curvelet transformation. *In Proc. International Conference on Information Assurance and Security* 2:215–218.