Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

On Soft Biometrics

Nixon, Mark; Correia, Paulo; Nasrollahi, Kamal; Moeslund, Thomas B.; Hadid, Abdenour; Tistarelli, Massimo

Published in:
Pattern Recognition Letters

DOI (link to publication from Publisher):
[10.1016/j.patrec.2015.08.006](https://doi.org/10.1016/j.patrec.2015.08.006)

Publication date:
2015

Document Version
Early version, also known as pre-print

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Nixon, M., Correia, P., Nasrollahi, K., Moeslund, T. B., Hadid, A., & Tistarelli, M. (2015). On Soft Biometrics. *Pattern Recognition Letters*, 68(Part 2), 218-230. <https://doi.org/10.1016/j.patrec.2015.08.006>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

On Soft Biometrics

Mark S. Nixon^a, Paulo L. Correia^b, Kamal Nasrollahi^c, Thomas B. Moeslund^c, Abdenour Hadid^d, Massimo Tistarelli^e

^a *University of Southampton, UK*

^b *Instituto de Telecomunicações, Instituto Superior Técnico, Lisboa, Portugal*

^c *Aalborg University, Denmark*

^d *University of Oulu, Finland*

^e *University of Sassari, Italy*

ABSTRACT

Innovation has formed much of the rich history in biometrics. The field of soft biometrics was originally aimed to augment the recognition process by fusion of metrics that were sufficient to discriminate populations rather than individuals. This was later refined to use measures that could be used to discriminate individuals, especially using descriptions that can be perceived using human vision and in surveillance imagery. A further branch of this new field concerns approaches to estimate soft biometrics, either using conventional biometrics approaches or just from images alone. These three strands combine to form what is now known as soft biometrics. We survey the achievements that have been made in recognition by and in estimation of these parameters, describing how these approaches can be used and where they might lead to. The approaches lead to a new type of recognition, and one similar to Bertillonage which is one of the earliest approaches to human identification. © 2012 Elsevier Ltd. All rights reserved.

1. Introduction

One of the oldest biometrics approaches derived from the work by Bertillon in the 19th century and one of the newest areas is that of soft biometrics. The work of Bertillon was superseded by fingerprints after showing early promise; the interest in soft biometrics is reflected in the increasing volume of papers that mention it. Unlike Bertillonage, soft biometrics is unlikely to be superseded as it can be used to reinforce biometric identification as well as be deployed alone in the analysis of data that conventional biometrics cannot handle or with invariant attributes that conventional biometrics cannot even approach.

Further, the UK Niche Records Management System has primary features such as a suspect's name, address and date of birth. Other desirable fields include factors used in Bertillon's approach such as gender, ethnicity and height. The FBI requires some demographic information to conduct a fingerprint-based background check. The information includes sex/gender, race, height, weight, eye and hair colour. These factors are often included on suspect identification forms, along with other descriptions of a suspect's build and clothing. These factors are primary human factors that we shall see can be identified from (colour) image data that are now termed soft biometrics.

As we shall find in this survey, there is a plethora of soft biometrics. The original formulation [Jain 2004a, Jain 2004b] concerned measures that can be used to aid recognition rather than for identification. The measures were suited to the discrimination between classes rather than individuals, and could be used to buttress recognition performance. Later techniques were to extend the definition, to concern estimation of personal characteristics, such as height and gender. Such factors have been of interest from the earliest days of biometrics, and are of interest given their ability to split populations. Given the pervading need for security in modern environments, there has been a concerted interest for the use of soft biometrics [Reid 2013a] largely to handle the low quality of video images where traditional biometrics cannot be applied. This is reflected in a theme of soft biometrics for identifying people, originally termed semantic biometrics [Samangooei 2008].

In this article we shall review the approaches that have been made in this new field. We shall provide an updated definition of the term soft biometrics and describe its advantages in biometric recognition. There is a contemporaneous review with similar aims [Dantcheva, 2015] and that offers an alternate view into this new and exciting topic. We shall concentrate on initial approaches and state-of-art techniques, aiming to establish the main basis of these new approaches and where they have developed to. We shall complete by suggesting new areas for development in this new and fascinating set of approaches.

1.1. Bertillonage

A variety of factors motivated a precise system for identification in 19th century France after the abandonment of branding criminals and a system of deportation. Limitations of the descriptions used to identify people, especially in size of forehead and colour of the eye motivated Bertillon to develop an anthropometric system to systematically describe people for identification based on their physiological traits in 1879 [Bertillon 1889]. Though the photographs were useful for confirmation of identity, the acquisition techniques were not standardised. A photograph is of use for verification of identity rather than to discover it, when relying on manual search. Bertillon noted the failings of the police identification and cataloguing system and developed his father's anthropological work to a more systematic method of identifying people. His system of anthropometrics, eponymously Bertillonage, outlined the tools and techniques for the careful measurement of:

- physical features including length/width of head, lengths of certain fingers and the dimensions of the feet, arm, right ear and standing height;
- descriptions of the dimensions of the nose, eye and hair colour; and
- the description and location of notable scars, tattoos and other marks.

The method for gathering these features was outlined in Bertillon's manual [Bertillon 1889] along with a set of diagrams (see Fig. 1). He implored 'police authorities...not to introduce special modifications of their own' for this could destroy uniformity, a practice echoed in modern biometrics databases and experimental practices. The measurements for a given individual were held on separate slides along with standardised photographs of the individual. The metrics of the system were chosen primarily to be simple so that they could be gathered accurately. The measurements were taken by a trained individual, though not necessarily a skilled one. Features were chosen to allow easy identification of points to begin and to end measurement. The success of Bertillonage came from its ability to reduce the probability of type 1 errors. Though two individuals may have very similar height, the chance of the same two having similar measurements for the other features is unlikely. Furthermore, Bertillonage inherently allowed for efficient discovery of an individual's existing measurement card and therefore their identity. Cards were stored according to specific range combinations of each metric in a given order. As such that once new measurements of an unidentified individual were taken then the identity of the individual could be easily ascertained.

Achieving great success and popularity in France, Bertillonage progressed to many countries by the late 19th century. Difficulties in cases such as West vs. West [Cole 2007] where Bertillonage appeared to be unable to distinguish people of similar appearance is often quoted as a reason for it being superseded by forms of identification such as fingerprint analysis (since the fingerprints of identical twins differ). It is of note that recently there has been a study on the potency of body measurements vs those of the face, for identification, in some ways redeeming Bertillon's original approach [Lucas 2015]. The study was based on using anthropometric measurements of 3982 individuals from the US Army Anthropometry Survey (ANSUR) database and concluded that "The body is more variable than the face and should be used in identification" also stating practical advantage.

These systems aim to reduce identity to a representative and measurable set of features, though not using descriptions of the human body as a whole. Measurements are taken in a controlled way, much the same as in modern biometrics, though lacking its sophisticated statistical and recognition techniques.

1.2. On the Development of Soft Biometrics Approaches

The earliest approach explicitly mentioning a form of soft biometrics appears to be by Wayman [Wayman 1997] who proposed the use of soft biometric traits like gender and age, to filter a large biometric database. An early motivation of the first soft biometrics was to augment conventional biometric signatures and envisaged that soft biometrics would be obtained separately, perhaps not from images originally used for recognition, and then used to enrich the biometric signature. A database was acquired that comprised of four fingerprint impressions of 160 users (obtained using a Veridicom sensor) of which a reduced set were assigned to known face images. The basic biometric features were fingerprint minutiae determined using a previously established technique and these were enriched by the soft biometric measures of age group, gender and ethnicity. (We shall be using gender and sex synonymously throughout this paper) A parallel approach had considered age group, gender and height [Jain 2004b] and as the subjects' height was unavailable, the replacement measures were selected from a random distribution. Gender and ethnicity information of users were automatically extracted from their face images using established techniques. These soft biometric measures were shown to vastly improve on standard fingerprint matching: e.g. for identification an increase from 86.4% to 90.2% at rank 1, with a similar effect in verification performance showing an increase of 4%

in the Genuine Acceptance Rate over a wide range of values of False Acceptance Rate. The improvement in performance is illustrated in Fig. 2 where the inclusion of soft biometrics clearly improves the fingerprint recognition rate.

The earliest definition of soft biometrics stated “soft biometric traits as characteristics that provide some information about the individual, but lack the distinctiveness and permanence to sufficiently differentiate any two individuals” [Jain 2004a]. An extended version of the definition of soft biometrics later stated “Soft biometric traits are physical, behavioural or adhered human characteristics, classifiable in pre-defined human compliant categories” later adding “In other words the soft biometric traits instances are created in a natural way, used by humans to distinguish their peers” [Dantcheva 2011b]. The latter definition is rather prolix and the former definition does not admit estimation of soft biometrics or human description. As such we shall define soft biometrics as: the estimation or use of personal characteristics describable by humans that can be used to aid or effect person recognition. We shall describe soft biometrics approaches with this new definition in mind.

The remaining strand of research in soft biometrics emerged to estimate the same parameters, but by using human vision. This accrues strong invariant advantages as a human can handle viewpoint and lighting change much better than computer vision techniques. As with gait biometrics, there is a unique advantage that the recognition approach can be used when images are acquired at a distance and other biometrics are at too low a resolution to be perceived. Even then, there is still information that can readily be perceived by human vision, yet is difficult to extract automatically. Samangooei examined how this information can be used to enrich the recognition process [Samangooei 2008]. They called these descriptions semantic annotations and investigated their use in biometric scenarios. They described a group of visual physical traits formulated and showed their use for recognition. As the visual attributes were derived from human labelers, greater care was taken to encompass known factors in psychology, such as memory effects, defaulting and owner variables. Some psychological aspects could not be considered, such as the cross race effect. Overall, the measures were demonstrated to have good recognition capability when used alone and proved an excellent addition to recognition capability, as consistent with the earliest form of soft biometrics.

We shall review these factors in the following sections for face and body and then for other soft biometrics. We shall show how the measures can be used to buttress recognition, to effect recognition, and can be estimated for the different biometric traits. The main soft biometrics that can be estimated for all modalities are age and gender, whilst some other soft biometrics such as weight and height can only be estimated from a single modality (in this case the body).

2. Soft Biometrics for Face

Undeniably, a significant number of soft biometric traits can be extracted from face images and facial movements. This generally includes gender recognition (i.e. man vs. woman), age categorization (e.g. child, youth, adult, middle-age and elderly) and ethnicity classification (e.g. Asian, Caucasian and African). These are often referred to as demographic traits and are very useful for more affective human-computer interaction (HCI) and smart environments in which the systems should adapt to the users whose behaviors and preferences are not only different at different ages but also specific to a given ethnicity and/or gender. Automatic demographic classification is also useful in many other applications such as content-based image and video retrieval, restricting access to certain areas based on gender and/or age, enhancing the performance of biometric identification systems, collecting demographic information in public places, counting the number of men or women entering a retail store and so on. Other soft biometric traits that

can be extracted from face images include Kinship information (i.e. verifying whether two persons are from the same family or not) and skin color.

2.1. Identifying faces by soft biometrics

The first approach to face verification using soft biometrics was described as using attributes that included face, age and gender [Kumar 2009]. This was motivated by the need to be able to recognize people in scenarios where pose and expression were not constrained. These were formulated using binary classifiers trained to recognize the presence or absence of describable aspects of visual appearance, thereby expressing gender as whether a subject was male or not, rather than male or female. The attributes and their classification for a face are illustrated in Fig. 3. This was accompanied by simile classifiers – that removed any manual labelling for training attribute classifiers instead using binary classifiers trained to recognize the similarity of faces, or regions of faces, to specific reference people. Application to the Labelled Faces in the Wild (LFW) database showed the promise of the new technique: an 80% true positive rate was achieved at 10% false positive rate, a much greater performance than contemporaneous techniques. The study also showed how people could outperform the new approach, suggesting room for further improvement.

A later study showed how comparative attributes could be used for identification. There, the attributes were derived by inviting people to label one subject compared with another [Reid 2013b]. These attributes included more detailed consideration of face components and 63 labelers derived the measures for 50 subjects. The approach achieved a 75% recognition accuracy at rank 1, rising to 100% at rank 6, that correct retrieval could be obtained with a small number of comparisons, and the labelers preferred the comparative labels to categorical versions that had been used previously.

More recently Klare et al. [Klare 2014] proposed a more elaborate approach, where human describable face attributes are exploited to perform face identification in criminal investigations. An automated attribute extraction algorithm to encode target repositories with the attribute information. The extracted attributes, such as eyebrows, chin and eyes shape are compared with the same attributes encoded in hand drawn police sketches. The experimental analysis demonstrates its applicability in forensic identification of suspects.

A recent study has considered the performance of humans in determining these attributes for faces [Han 2014]. The approach extracted previously proposed biologically inspired features from face images and selected soft biometric features using a boosting algorithm. The performance was compared with human labels derived by crowd sourcing. The approaches were applied to several popular databases including a subset of LFW and a large proprietary database. The results show that humans can discriminate gender slightly better than the automated techniques (e.g. for male subjects a 99.6% Correct Classification Rate (CCR) by humans opposed to 98.7% for automated techniques) whereas automated techniques generally outperform humans at age estimation (e.g. a 4.5 years Mean Absolute Error (MAE) by humans opposed to 3.8 years for automated techniques since human appeared to overestimate the ages of individual subjects). This not only enables better appreciation of the attributes used previously, but also can be used to calibrate the estimates of soft biometrics to be discussed later. One new approach on soft biometrics rather than on attributes has found modest performance advantage when soft biometrics were used to augment traditional biometrics approaches of Local Region PCA and Cohort Linear Discriminant Analysis [Zhang 2015]. Via analysis on the Good, the Bad and the Ugly (GBU) Dataset the slight advantage was perhaps since false matches are more likely between faces of people sharing the same soft biometric traits. Another novel approach, originally proposed to recognize a wide range of facial attributes, has been evaluated in the recognition of facial

expressions, gender, race, disguise and beard [Mery 2015]. The performance achieved is at least comparable with the state-of-the-art techniques and with a wider basis.

2.2. Exploiting characteristic features in face images

Human faces always contain characteristic patterns which may alone provide a support for classification. Early works, particularly in the area of biological vision [Koch 1985, Itti 1998] aimed to explain the role of fixation in distinguishing subjects. This line of research, coupled with findings in the area of forensic face analysis [Spaun 2007], lead to a better understanding on how to extract and exploit distinctive facial patterns such as scars, moles, tattoos, as well as any other peculiar feature, for face identification.

[Bicego 2008] proposed to compute the most distinctive facial regions by comparing a face image with face images from other individuals. The algorithm is based on the computation of face differences to determine the level of distinctiveness of any given face image. Space-variant patterns are randomly sampled from the face image, obtaining a large number of scale-invariant local features. A feature was selected as distinctive if it was significantly different from any other feature in a given set (like a facial scar or a mole). A perceptual experiment involving 45 observers, performed on the Banca database, indicated the output of the algorithm to be fairly comparable with how humans mark distinctive facial features. In [Park 2010] a statistical classification of image patterns into facial marks is performed. Facial marks are detected by means of a blob extractor and subsequently classified. The relevance of the texture color of facial marks is also addressed. For example, a mole on the face is generally darker than a spot, and a scar on the face is a discolored region. Color and shape information are combined to categorize the facial marks into different semantic categories (see Fig 4). Experimental results on the FERET and two forensics mugshot databases demonstrate the usefulness of facial marks to enhance the performance of commercial face recognition systems.

Several authors proposed to exploit texture patterns extracted from the periocular region of a face image as a soft biometric trait [Lyle 2012, Merkow 2010, Park 2011]. [Park 2009] firstly established the utility of periocular images as a soft biometric trait for classification purposes, while [Lyle 2012] investigated the effectiveness of local appearance features, such as Local Binary Patterns, Histograms of Oriented Gradient, Discrete Cosine Transform, and Local Color Histograms, extracted from periocular region images, for soft classification on gender and ethnicity. In [Park 2011] an extended analysis of the usefulness of periocular distinctive patterns was performed. A number of issues are considered, including the relevance of different parts of the periocular region and the effects of aging, disguise and occlusion. Experimental results show a rank-one recognition accuracy of 87.32% using 1136 probe and 1136 gallery periocular images taken from 568 different subjects (2 images/subject) in the Face Recognition Grand Challenge (version 2.0) database with the fusion of three different matchers. A comparison between face recognition and periocular recognition performance under simulated non-ideal conditions (occlusion) was also presented.

2.3. Gender classification from face images

The first attempts to use computer vision based techniques for gender classification started in early 1990s [Golomb 1990]. Since then, significant progress has been made and several approaches have been reported in literature. Fundamentally, the proposed techniques differ in (i) the choice of the facial representation, ranging from the use of simple raw pixels to more complex features such as Gabor responses, and in (ii) the design of the classifier, ranging from the use of nearest neighbor (NN) and Fisher linear discriminant (FLD) classifiers to artificial neural networks

(ANN), support vector machines (SVM) and boosting schemes. For instance, Moghaddam and Yang [Moghaddam 2002] used raw pixels as inputs to SVMs while Baluja and Rowley [Baluja 2007] adopted AdaBoost to combine weak classifiers, constructed using simple pixel comparisons, into single strong classifier. Both systems showed good classification rates of above 90% but under controlled settings. An empirical comparative analysis on gender classification approaches can be found in [Makinen 2008] whereas a recent survey on the topic can be found in [Ng 2012]. There has also been some concern that cosmetics can affect the perception of gender (and age) [Chen 2014]. Naturally cosmetics and plastic surgery can affect soft biometric estimation and we anticipate future work in these areas (in line with works in biometric antispoofting). Note that there is no public database specifically designed for gender recognition evaluation. Most of the recent works have been evaluated on the Images of Groups database [Gallagher 2009] which is a collection of face images of groups of people from Flickr. This database has been mainly designed for unconstrained demographic classification (age and gender).

2.4. Age classification from facial images

Automatic age classification aims to assign a label to a face regarding the exact age (age estimation) or the age category (age classification) it belongs to. This is challenging problem because the appearance of a particular face varies due to changes in pose, expressions, illumination, and other factors such as make-up, occlusions, image degradations caused by blur and noise etc. In addition to these difficulties which are shared with the standard problems of face recognition, ageing is a very complex process that is extremely difficult to model: a group of people of the same age may look very different depending on, for example, environment, lifestyle, genes etc. Thus, deriving a universal age classification model is troublesome. Many face image representations for age estimation have been studied such as anthropometric models, active appearance models (AAM), aging pattern subspace, and age manifold. An extensive review of age representation methods can be found in [Fu 2010]. Regarding age classification schemes, the existing methods are based on either pure classification or regression analysis. Perhaps, among the pioneering studies on age classification are those proposed by Kwon and Lobo [Kwon 1994], Lanitis et al. [Lanitis 2002], and Guo et al. [Guo 2009a]. More recent approaches have considered short-term and long-term aging as different phenomena [Ortega 2009, Yadav 2013], where the estimation framework is based on the direct measurement of the time evolution of the face shape. Although relatively successful in some scenarios (e.g. high-quality and occlusion-free images, neutral facial expression), most existing methods tend to suffer under uncontrolled settings as noted in [Fu 2010]. The most commonly used public face databases for evaluating age classification are MORPH [Ricanek 2006] and FG-NET [Panis 2014] database and the more challenging Images of Groups [Gallagher 2009] database.

2.5. Race and ethnicity classification from facial images

While gender and age recognition have been explored by many other researchers, automatic ethnicity classification problem has received relatively far less attention despite the potential applications. This is perhaps due to the ambiguity and complexity in defining and describing different ethnic groups. The terms “race” and “ethnicity” are sometimes used interchangeably although they refer to biological and sociological factors respectively. Generally, race refers to a person’s physical appearance or characteristics, while ethnicity is more viewed as a culture concept, relating to nationality, rituals and cultural heritages, or even ideology [Fu 2014]. The most commonly encountered and accepted racial groups are African/African American, Caucasian, East Asian, Native American/American Indian, Pacific Islander, Asian Indian, and Hispanic/Latino. A recent survey on extracting race from face images can be

found in [Fu 2014]. Most of the general face databases are not intentionally designed for race or ethnicity classification. Another very recent survey [Arigbabu 2015] on soft biometrics for face recognition concentrates on feature extraction and application of soft biometrics.

3. Soft Biometrics for Body

3.1. Describing the whole body for recognition

The first approach to soft biometrics based on human description [Samangooei 2008] obtained soft biometric labels for each subject in the Southampton gait database [Shutler 2002]. The labels were collected via a web interface and known psychology factors were taken into account. These factors included:

- Memory – the labellers were allowed to view images indefinitely so that memory could not impair the labels provided;
- Defaulting – the labellers were not provided with default labels, but were explicitly required to state values for labels;
- Anchoring – which could occur from the phrasing of the label was addressed by using “unsure” as opposed to “average”. Anchoring could also occur due to the order in which subjects were presented and so this order was randomised;
- Categorisation – labellers were provided with five distinct categories for each label; and
- Owner variables – since labels are likely to be influenced by a labeller’s perception of themselves, their description of themselves was also collected.

The only known and pertinent psychology factor that was omitted was the cross race effect, since ethnicity is notoriously unstable. It is hard to imagine how a factor which more affects recognition than a labelling structure could actually be accommodated in this scenario. Each of 10 subjects selected from the Southampton gait database had annotations from 38 annotators using the web interface illustrated in Figure 5.

The labels themselves were directed by an earlier study from psychology wherein the body traits used by humans to identify each other were studied [Macleod 1994]. There, labellers watched video footage of subjects walking at a regular pace around a room and rated them using 23 traits identified from human descriptions of physique and motion. Statistical analysis of the 23 traits led to 13 that were the most significant (as the most representative of their principal components). The traits chosen were suitable for deployment at a distance allowing also for analysis of surveillance data and not restricting the approach.

Some examples of the 23 traits and their descriptions (terms) are given in Table 1. These show that the traits describe body and face features. Traits were usually described using a five point scale, though the age appeared to need a much larger set of descriptions particularly for younger subjects and this reflects the rate of change in appearance in youth. The face features collected were those available at a distance, rather than the fine grained identifications used for facial image identification. These descriptions of were accompanied by descriptions of age, ethnicity and sex.

An investigation of the correlations between these semantic labels showed how often individual trait and term pairings were used by annotators [Samangooei 2014]. The more informative correlations were observed between traits whose terms describe overall thickness and length of the body, as well as extremities. As such, Figure and Weight were highly correlated, and in turn they are both correlated with Arm Thickness, Leg Thickness and Chest annotations. Correlation was also noted between Height and Leg Length, each also portraying correlations with Arm Length. There were some inverse correlations such as between arm and leg shape and shoulder shape and many other

measures (though this is perhaps more due to the inability to observe shoulder shape from a side view). The Race and Sex features were found to be statistically the most significant (in terms of p-value) of all the descriptions used.

When the descriptions were deployed for identification, 90% CCR could be achieved when used alone and more than 40% of the features were incorporated in the feature vectors. Score fusion was then used to improve the recognition achieved by automatic gait biometrics: when an established discrete-symmetry based gait biometric was deployed alone a 98.1% CCR was achieved and with score fusion of the semantic features this rose to 99.5% CCR. Clearly the semantic features could be used alone, or to buttress conventional biometric approaches.

An extension to Samangooei's soft biometric system was motivated by an analysis of the subjects' perception of height. In Samangooei's study, this had been labelled by users in a categorical way wherein the labels depended on the labellers and on their impression of scene geometry (only the positions of the ceiling and the lamp give height information in Fig. 5). It was observed that a subject who was in reality 310 pixels tall could be given labels of Short, Medium or Tall. The labels Very Tall and Very Short were rarely confused. This was solved by changing the structure of the semantic terms to be relative rather than categorical: the labels were derived by comparing one subject with another using a new Web interface. For example the Height trait was labelled as Much Shorter, Shorter, Same, Taller or Much Taller. The traits Gender, Ethnicity and Skin Colour remained categorical, whereas Age was described by a reduced set of terms, ranging from Much Younger to Much Older. Again, 'politically correct' labels were avoided to ensure labellers understood the precise targets: Weight was now labelled as Thinner or Fatter; estimating Height avoided any concept of vertical challenge. 19 comparative traits of 100 subjects from the Southampton gait database were labelled by 57 annotators, together with three categorical traits.

Since the labels were now comparative, they needed to be sorted into their categorical version for use in recognition. This was achieved using the Elo ranking system, though ranking SVMs could have achieved a similar purpose. After sorting, the labels on Height were observed to have a linear relationship with Height measured in pixels, with correlation 0.87. When used for recognition, 95% CCR was observed with 20 comparisons per subject and 85% CCR with only 5 comparisons. As such the relative description process considerably improved recognition capability via semantic labels. A later extension of the approach predicted the labels using gait biometrics approaches, and although the retrieval accuracy was only 20% at rank 1, it quickly increased to 69% at rank 9 and a 90% at rank 19 [Reid 2014]. As such, it could then be possible to automatically determine the identity of subjects within video footage by using verbal descriptions, and research on these approaches continues. New work shows that many and accurate labels can be achieved by a crowdsourcing system [Martinho 2015].

3.2. Estimating gender from images of the whole body

In contrast with the high volume of work aimed to estimate gender from face images, the estimation of gender from other data has received much less attention. There have been works aimed to estimate gender automatically from static images of the whole human body [Cao 2008, Collins 2009] and works aimed at deploying gait biometrics for gender estimation. The first approach used full body images, taken from front or back views. Each image was transformed into a collection of patch features which modelled different body parts to provide clues for gender recognition. The transform was based on the use of the Canny edge operator and the Histogram of Oriented Gradients (HOG), classified using Adaboost and Random Forests. The results achieved at most a CCR for gender of 76% for images from the front view (suggested considerable room for improvement given that the random rate for a balanced dataset is around 50%). A contemporaneous approach, largely based on machine learning rather than computer vision, achieved around 80% successful classification of gender from images of the whole body [Guo 2009b]. A

slightly later approach based on techniques including pyramid HOG and spatial pyramid bag of words slightly improved on performance [Collins 2009], noting that “a publicly available database of human body pictures with gender labels does not currently exist”. A more recent approach has concerned estimating gender and weight from geometric measurements of human subjects [Cao 2012]. These measures were derived from the head and from the remaining parts of the body. The model resulted in a 0.7% misclassification rate for gender prediction using both body and head information, 1.0% using only body information, and 12.2% using only head information using a database of over 2000 subjects. The approach clearly suggests that human body metrology contains sufficient information to reliably predict gender, and appears to give the most accurate performance to date for estimating gender. The earliest approach to determining gender from gait (and thus from sequences of images) showed considerably greater discriminatory capability achieving 96% CCR on sequences of side view images [Yoo 2005] inspiring much later work since knowing the gender of a suspect in a criminal proceeding could eliminate half the population from enquiries. This result was perhaps due to the laboratory nature of the data which certainly confirms that such a task is possible. A later approach which concerned video rate extraction of soft biometrics observed a lower rate of 79% CCR when applied to the USF gait data which was derived outdoors in a much less constrained environment [Ran 2008]. In the only approach to use semantic labels to retrieve images of subjects of a specific gender [Samangooei 2009] used Latent Semantic Analysis with a dataset of 2000 videos of people walking in laboratory conditions and achieved promising retrieval results for features such as Sex (mAP = 14% above random) and Age and Ethnicity.

3.3. Estimating subjects' height from whole body images

There has been a long history in determining suspects' height in video, not just for identification but also since height is a strong indicator of a subject's gender. There have been forensics approaches, based on photogrammetry [De Angelis 2007] and on computer vision [Criminisi 1999]. As already mentioned, an early gait biometrics approach [BenAbdelkader 2002] included automatic height estimation. A later approach used gait to direct height estimation by focusing on subject height by estimating the body frame size from the silhouette corresponding to the double support gait pose when a human is at maximum height (the human appears shortest at heel strike which is when the leading foot first strikes the ground), achieving between 75% and 87% CCR (with a potential 10% error rate) [Ran 2008]. A slightly later approach [Denman 2009] also used human height, estimated by statistical means and observed difficulty with errors in feet localisation and similarity of subject height in part of the PETS dataset. It is worth noting that the earlier semantic approaches were also interested in human perception of height, but this was not correlated with ground truth (where available) since the intention was to explore the ability of human descriptions for recognition, rather than human descriptions for estimating soft biometrics parameters.

3.4. Estimating subjects' weight from whole body images

One study has concerned human ability to estimate subjects' weight from whole body images [Velardo 2010]. The study suggested a set of measures that comprehensively covered the whole body (upper and lower part) and were reasonably correlated to the weight. The measures included height, upper leg length, calf circumference, upper arm length, upper arm circumference, waist, upper leg circumference. The approach used manual estimates from a nutrition survey and showed that the circumference of the upper arm was most correlated to weight. By using manual annotation of front-view and side-view images of 20 subjects the system was able to estimate the weight from visual clues with an approximation of $\pm 5\%$ of error over the real weight of the subject. As such, the analysis confirms the human ability to estimate weight and suggests a possible approach to future automation. The approach mentioned

earlier for estimating gender and weight [Cao 2012] gave 0.01 mean absolute error for weight prediction using the body and rather larger error (0.07) using the head. The approach clearly suggests that human body metrology contains sufficient information to reliably predict gender and weight.

3.5. Estimating subjects' age from images of the whole body

To our knowledge the only approach aimed to estimate a subject's age from a single image of their body was a feature based approach that extracted the SIFT feature and then applied sparse coding to learn a dictionary for feature quantization [Ge 2013]. Linear regression was then used to learn the relationship between age and the feature representations. On a database of WWW images of 1500 subjects, ranging from very young to very old and largely of young to middle age, the Mean Average Error was 8.76 years. The performance was less than that observed for age estimation from images of the face [Guo 2009] where the MAE was 4.7 years. The first approach to estimating age from gait biometrics derived features from the gait energy images [Lu 2010]. The approach was evaluated on the USF database which contained around 100 subjects, largely in their 20s, achieving a MAE of 6.23 years. A following approach [Makihara 2011] used the gait energy image within its feature set and used Gaussian Process Regression to determine the age. This was applied to a database of nearly two thousand subjects, ranging from young to very old, and achieved an MAE of 8.2 years which is close to that observed when using single images rather than sequences. It will be interesting in future to note the relative advantages associated with sequence- or single-image-based techniques.

3.6. Exploiting correlation to predict body characteristics

As in any exercise in pattern recognition it is possible to explore the nature of the feature descriptions to improve retrieval and recognition performance. An investigation into the structure of the labels has already been described [Samangooei 2014] which exposed expected correlations. These have previously been used to impute or predict the values of labels that have been omitted or could not be observed [Adjero 2010, Reid 2010]. One approach used continuous data focusing on measurements of the human body [Adjero 2010]. Data was gathered from the CAESAR anthropometric dataset which comprised of 45 human measurements or attributes for 2369 subjects and assessed relationships between the human measurements using the Pearson correlation coefficient. A correlation graph illustrated connections between traits where the correlation exceeded a threshold value. It was observed that the measurements generally fall into two groups (both with physical meaning): a 2D group containing body parts' perimeters and a 1D group measuring length (and height). The clusters suggested that only a few measurements would have to be known to predict the majority of the other traits. Four 'seed' measurements (arm length, knee height, shoulder breadth, and standing height) were used to predict other traits with low average mean absolute error. The study also showed how weight and gender could be estimated from the measurements, achieving 88.9% CCR for a testing set.

An alternative approach concerned study of correlation for prediction from the human labels [Reid 2010]. This was aimed to handle known performance restrictions when labelling surveillance video, such as occlusion of the body either by position or by gesture. Entropy was used as a measure of certainty as to how successful a term is for predicting a missing trait. It was observed that the traits most successfully predicted were skin colour and ethnicity, and that this was likely due to strong correlation with other traits allowing accurate prediction of missing data.

4. Other Soft Biometrics

Soft biometric information is often computed from face or body as discussed in previous sections. However, additional information can be considered, notably related to scars, marks and tattoos (SMT) [Lee 2008] or related to accessories, such as the presence of eye glasses or clothing characteristics, even if in such cases the two factors of *distinctiveness* and *permanence* assume a particular importance [Dantcheva 2011b]. For instance, an application that performs identification of subjects within one day or re-identification across cameras, may consider clothing colour as having the desired permanence, while in general the permanence of clothing traits would be considered low. This section focuses on the usage of clothing attributes, as well as colour as soft biometrics. Also gender estimation from traits other than face and body is discussed – examples are included for gender from iris, fingerprints or ear. Similar examples can be found in the literature for the estimation of age using alternative features, as briefly discussed in the following sections.

4.1. Clothing

A systematic approach to the description of people by their clothing properties is presented in [Jaha 2014]. It corresponds to a task naturally undertaken by humans, notably to identify/re-identify each other, especially from a distance, fitting our definition of soft biometrics. Fig. 6 illustrates a situation where clothing attributes were visible and can be used as a soft biometric, contributing to the identification of a suspect.

In 2012 Chen et al. [Chen 2012] proposed an automated system for generating a list of nameable attributes for human upper body clothes, in unconstrained images. The semantic attributes considered include: Clothing Pattern (solid, floral, spotted, plaid, striped, graphics), Major Colour, Wearing Necktie, Collar Presence, Gender, Wearing Scarf, Skin Exposure, Placket Presence, Sleeve Length, Neckline Shape, and Clothing Category (shirt, sweater, T-shirt, outdoor, suit, tank top, dress). The system starts by estimating the human pose prior to feature extraction and

performs attribute classification using SVMs. Reported results support the importance of considering clothing attributes for identification purposes.

Jaha et al. propose a more complete description of semantic clothing attributes, covering the full body, notably: Head, Upper Body, Lower Body, Foot, Attached to Body, General Style, and Permanent, as detailed in Table 2. Reported results support that clothing soft biometrics show potential for identification even when used alone, and that their combination with traditional soft biometrics allows a substantial improvement of the otherwise obtained results.

Recently a number of studies have addressed the issue of person re-identification across non-overlapping camera views. For this purpose clothing attributes are often the main available cue, as the appearance of people is most influenced by their clothes, which assumes an especially important role for footage captured at a distance. In [Li 2014], a representation of several overlapping body parts is considered, with colour features being computed for each. Then middle-level clothing attributes are used to bridge the semantic gap between low-level features and the desired high level classification. The considered clothing attributes are grouped into the following categories: Head (bold, hat, short hair, long hair), Shoulder (epaulet, no epaulet), Sleeve (long sleeve, short sleeve, no sleeve), Shoe (multiple colour, single colour), Style (dress, skirt, shorts, longs single colour, longs multi-colour), Apron (apron, no apron), Open Coat (open coat, no open coat), Front Pattern (front pattern, no front pattern), Back Pattern (back pattern, no back pattern), Carrying (backpack, satchel, handbag, tray, no carrying), and Texture (upper texture, lower texture, whole texture, bag texture, no texture). Clothing attributes can be annotated as discrete values, or relaxed to continuous values to reflect the confidence of assigning a given value. A latent SVM framework is employed by [Li 2014] and tested on a database collected from 10 cameras, with annotations available for a total of 215 persons. Results confirm that adding clothing attributes considerably improves re-identification results. The earlier clothing attributes approach has been extended for re-identification [Jaha 2015] since most biometrics are not visible when comparing front to side viewpoints. As clothing is visible from both viewpoints, and can be used for identification, it gives an effective means for re-identification. A new approach via clothing attributes is part of an “actual product for people search in surveillance videos” where the ‘fine-grained’ clothing attributes include color, clothing type and pattern [Chen 2015].

Following a similar reasoning [Layne 2012] analysed the operating principles of human experts to propose an ontology of attributes for re-identification purposes. The 15 binary attributes considered are: Shorts, Skirt, Sandals, Backpack, Jeans, Logo, V-Neck, Open-Outwear, Stripes, Sunglasses, Headphones, Long-Hair, Short-Hair, Gender, and Carrying-Object. Among these we find 3 of the traditional soft biometrics, which are combined with 12 additional soft clothing traits. Attribute detection starts with low-level feature extraction using colour and texture features, fed to a set of SVMs that are combined with spatial constraints to detect the desired attributes.

4.2. *Colour*

Colour information extracted from different sources, as mentioned before, has often been used as a soft biometric. Such sources include: clothing [Møgelmoose 2008], skin [Jaha 2014, Park 2010, Wallace 2013], hair [Dantcheva 2010], and eyes [Dantcheva 2011a].

Using colour of clothing in surveillance videos has mainly been used in systems where tracking of humans within one (or multiple) camera’s point of view [Okuma 2004] or re-identification from one camera to another [Møgelmoose 2008] is of interest. In such systems usually the human body is divided into two or three parts and then these parts are modelled by different techniques like mixture of Gaussians or histogram [Hansen 2007]. In [Møgelmoose 2008] an RGB-D system has been developed for person re-identification in which having segmented the person, the histogram

of the colour of his/her clothing has been used as a soft biometrics. To do that, as shown in Fig 7, the body is first divided into upper and lower parts and then a separate colour histogram is calculated for each body part. Each histogram contains 60 bins, in which 20 bins are assigned to each of the R, G, and B channels.

In [Dangelo 2010] colour of clothing has been used as a soft biometric for identifying rival team's fans in surveillance videos for the purpose of preventing fights. To do so, they used a fuzzy k-nearest neighbour method for classifying of colour pixels in surveillance videos to a so called set of cultural colours composing of 11 colour bins which has been learned by colour samples of the clothing of some sport teams from different cameras [Dangelo 2010].

Skin colour, along with age, sex, and ethnicity, has been one of the traditional soft traits [Jaha 2014]. In [Samangoeei 2008] skin colour has been semantically divided into four semantic groups of white, tanned, oriental, and black. As discussed in [Wallace 2013], complexion features (like skin colour and texture) provide distinctive information similar to features like shape of face, for humans. In [Wallace 2013] a histogram based method using facial skin and texture has been developed in HSV colour space (using only H and S components), and the importance of these features in low-resolution facial images has been studied. The robustness of the features against changes in the camera's white balance has also been reported. Similar discussion under different imaging parameters for eye colour has been reported in [Dantcheva 2011a]. In [Park 2010] the importance of the colour of facial marks as soft biometrics is studied. They have provided a classification for the different types of facial marks, using different factors, among them colour. They discuss that, for example, a mole on the face is darker than a spot, and a scar on the face is a discoloured region. They generally classify the facial marks into two groups of dark or bright compared to their surrounding facial regions. This colour information is then combined with the shape of the detected marks to categorize them into different semantic categories. Experimental results on FERET database and two forensics mugshot databases in [Park 2010] show improvement in the performance of a face recognition system which uses these facial marks.

4.3. Gender from Other Biometrics

Gender is one of the most widely used soft-biometrics. Besides the usual techniques for gender recognition based on facial and body information, additional traits can be considered. A few examples are described in this section. Gender classification from iris images has been addressed for instance in [Thomas 2007, Lagree 2011]. The first approach used texture and geometric features aiming to derive "a feature vector from iris images to predict gender". The accuracy by a variety of approaches neared 80% CCR on a database of 57,137 images (aimed to balance between male and female) derived using an Iridian LG EOU200 system. The later approach predicted both ethnicity and gender (improving the classification capability for gender from iris), and studied the mixed effects of the two problems.

Fingerprints can also be used to estimate gender as can speech [Li et al 2013]. The analysis of several texture descriptors to estimate gender from fingerprint images is presented in [Rattani 2014]. The considered descriptors are: LBP, local phase quantization (LPQ), binarized statistical image features (BSIF) and local ternary pattern (LTP). Combined with an SVM classifier LBP, BSIF and LTP are able to achieve a CCR of around 71% on the West Virginia University (WVU) database. If majority voting fusion for the subject's four available fingerprints is considered, results increase to around 80% for the same descriptors. When testing the ability to estimate gender using fingerprints originating from different fingers, a reduction in performance is observed, as expected, with LBP results achieving 66%. [Rattani 2014] also studies the robustness of texture descriptors to several types of image

degradations and to the usage of incomplete fingerprints, as one may encounter in forensic scenarios. In the presence of noise the BSIF descriptor is the more robust, with a CCR of 52.6%. For partial palmprints the BSIF descriptor is little affected, again showing the best performance. On the other hand, LPQ is quite robust to blur, showing only a small performance decrease.

Since 2D images can change considerably with pose variation, it has been argued that there is advantage in also considering 3D information obtained from range data. One such example is the usage of 3D ear shape, as ear contains little texture and the most descriptive attribute for gender recognition is its shape [Lei 2013]. The paper starts by the detection of a set of ear landmarks, followed by normalizing of 3D ear images to reduce the effect of pose variation, and then tests recognition effectiveness using three features: ear curvedness map (ECM), histogram of indexed shapes (HIS) and surface patch histogram of indexed shapes (SPHIS). Neural networks, SVM and Adaboost classifiers were tested for estimation of subject's gender. Tests were conducted using the University of Notre Dame datasets F and J2, containing 942 3D face profile scans, and 1800 3D profile face range images, respectively. The test data included scale and pose variations as well as occlusions by hair and ear rings. The best results were achieved for the HIS feature using a SVM classifier, with an accuracy rate of 92.4%. Moreover, [Lei 2013] shows better results are achievable with 3D over 2D data.

4.4. Age from Other Biometrics

Age is a soft biometric that can be not easy to estimate, especially for younger people. A recent book discussing age factors in biometrics processing addresses many of the aspects related to ageing and how they are reflected in biometric traits [Fairhurst 2013]. Some examples based on traits other than face and body have been considered, such as iris or handwritten signature, as well as their combination [Sgroi 2013, Erbilek 2014]. Other examples include age estimation from voice [Bahari 2011, Li 2012, Mendonza 2014], fingerprints or palmprints [Lanitis 2011].

5. Future Work

There are many areas innate to pattern recognition and to biometrics that are of interest to the advance of soft biometrics. Feature set selection and machine learning for recognition are of natural interest here, as is the study of experimental factors/ covariates, and of fusion of soft biometric traits. Equally, the application areas of biometrics can benefit from soft biometrics. There are new moves in the presentation of forensic evidence on a probabilistic basis, for which biometrics and its databases appears most suited. This applies equally to soft biometrics, though there has as yet been no focused study on the use of soft biometrics in forensics.

Though there has been a great deal of progress in face analysis in the last two decades, most work has mainly focused on face detection and recognition, rather than on soft biometrics. Consequently, the design of algorithms that are effective in discriminating between males and females, or classifying faces into different age and ethnic categories remain open areas of research. Findings from human perception studies, suggested that peculiar texture patterns, sparsely located on the face image, are often used by humans for a rapid identification of faces. The same approach demonstrated to be applicable in forensic cases, as well as other applications, as a soft biometric trait. Such isolated patterns, which may correspond either to a skin defect (such as a scar or a mole), or to a specific facial area (eyebrow, chin) with a distinctive appearance, often do not allow for an accurate classification but can be exploited in conjunction with other techniques to improve the recognition performance. Some of these traits have been

successfully applied for age, ethnicity, and gender estimation. The difficulty with race and ethnicity classification might prove to be an interesting area of future work, perhaps clarifying not only terminology but also capability.

While many works on demographic classification reported accuracy of above 90%, most of these studies have utilized face images captured in rather controlled sensing and cooperative subject scenarios [Han 2014]. However in many real-world applications such as video surveillance, the available face images of a person of interest are most likely to be captured under unconstrained and un-cooperative scenarios. Most of the proposed techniques for demographic classification have been evaluated on general face databases which are not intentionally designed for the problem at hand. Thus, there is clear lack of benchmark databases and protocols to evaluate the progress in extracting soft biometric traits from face images. Recently, efforts have been made within the BeFIT project (<http://fipa.cs.kit.edu/befit/>) by proposing standardized datasets and protocols for evaluating different face analysis tasks. We strongly endorse this valuable initiative which allows for fair comparison and easy reproduction of research results. This can apply to all areas of interest in soft biometrics, though some such as face are much more suited to situations where the subject is imaged in great detail, and body where the resolution is much less.

The way a person moves his/her head and facial parts (such as the movements of the mouth when a person is talking) defines so called facial dynamics and characterizes personal behaviors. An emerging direction in face analysis consists of also using such dynamic cues, in addition to facial structure, in order to enhance the performance of static image based methods. An example reporting the exploitation of facial behavior is available at [Kashyap 2012]. However, current works on extracting soft biometric traits are mostly focused on static 2D face images while ignoring facial dynamics [Hadid 2013]. This applies to body dynamics too, and there has been resurgent interest in identification by gesture. At the moment this is either stylized behavior patterns (such as sitting and jumping) or stylized surveillance requirements (such as leaving a bag) [Denman 2012]. An example analyzing 8 simple hand gestures, for a small database of 20 people, is reported in [Lai 2012], achieving an EER reduction of 8-10% in comparison to single silhouette based recognition. Other examples considering handwriting motion in the air [Takeuchi 2013] or authentication based on sign language gesture images [Fong 2013] have been reported. The identification by gesture could then consider the body soft biometrics already developed and extend these in a spatiotemporal manner.

Gesture can also be analyzed for biometric recognition when interacting with mobile devices. There has been a considerable amount of work on using keystroke dynamics for recognition [Morales 2014, Idrus 2014], and more recently also touch gesture analysis is being considered for recognition purposes [Meng 2012, Ruebsamen 2013]. Additionally the concept of social behaviour biometrics has been proposed embracing the usage of soft biometrics [Sultana 2014].

2D based demographic face classification methods usually encounter the difficulty when facing geometric and illumination variation, to which 3D model based approach is rather insensitive. So, it is of great interest to investigate the extraction of facial soft biometrics from 3D images or even in 4D (i.e. 3D in videos). This applies equally to body and other soft biometrics. Certainly one advantage of 3D information is that some information is available in one view and not in another. This can allow for selection of an appropriate viewpoint selection for comparison with a particular sample. This generalizes to the problem of re-identification wherein subjects might reappear in surveillance video in a different pose to their previous version. Also the emergence of high dynamic range (HDR) cameras may have a considerable impact in the ability to handle otherwise poorly illuminated scenes, enabling a better extraction of soft biometric information even in more un-constrained environments.

A primary advantage of soft biometrics is their potential for deployment at a distance. As yet, only one study has focused on distance and soft biometrics [Tome 2014]. This is an interesting area of study since it is known that the suitability of a chosen modality depends largely on distance, and experimental study of the effect of distance can help to guide selection. Age, gender and ethnicity are naturally inter-related and hence should be studied in a joint framework [Farinella 2013]. For instance, it is known that it is more difficult to determine the gender of a child compared to that of an adult. Another potential area of deployment is in targeted marketing, since soft biometrics include factors of primary interest, especially gender and age.

Given that the technology is likely to embrace surveillance, and the growing concern on deployment of surveillance technology, it is likely that there will be ethical considerations that affect not just biometrics [Mordini 2008, Kindt 2013], but also soft biometrics. Of particular interest is ethnicity which becomes less stable in a modern multicultural society. There are also the privacy issues that are of concern to any biometric, especially to those ripe for deployment in surveillance scenarios.

Primary motivations of biometrics include security and convenience: a walk-through biometric scanner used in immigration can be viewed as epitomising progress in automatic identification. The advantage of soft biometrics is that they can improve both factors. Soft biometrics share factors with many biometrics such as the possibility of spoofing mandating antispooofing techniques, the need for fusion, the depth and types of features used require further study. The automatic estimation of soft biometrics could aid the retail market. This could be achieved by targeted advertising (highlighting objects that are likely to be of interest to a viewed subject) and by marketing (confirmation of age by soft biometrics could save time in, say, the purchase of items that are restricted to older people. The final strand of soft biometrics is identification in surveillance video and this reflects a wider interest in image-based search. As these systems are now starting to emerge, it appears that soft biometrics is a rich area for future work.

6. Conclusions

There has been much progress in soft biometrics. There have been several definitions and to encompass these and to set some clear boundaries for future study, our new definition is that soft biometrics are: the estimation or use of personal characteristics describable by humans that can be used to aid or effect person recognition. Their original formulation concerned buttressing biometrics using broad measures. At the same time there has been much work on estimating soft biometrics, especially from face images. This is natural, since face recognition has a history almost as long as fingerprint recognition and is of continuing interest in society. The remaining strand of research concerns identification by soft biometrics, especially by human descriptions. Each of these formulations has been applied to face images, to images of the human body and to other biometrics. Soft biometrics have been demonstrated to enrich biometric performance, to be able to be deployed for recognition and to be used for estimation of factors that are of interest in many applications. We show how this has been achieved, the current state of art, and where this might lead.

Acknowledgments

This work has been partially supported by the European Commission COST Action IC1106 "Integrating Biometrics and Forensics for the Digital Age"

References

- [Adjeroh 2010] D Adjeroh, D Cao, M Piccirilli, A Ross, Predictability and correlation in human metrology, Proc. IEEE Workshop Inf., Forensics and Security, 2010
- [Arigbabu 2015] OA Arigbabu, SMS Ahmad, WAW Adnan, S Yussof Recent advances in facial soft biometrics, The Visual Computer, 31:513-525, 2015
- [Bahari 2011] MH Bahari, H Van Hamme, Speaker age estimation and gender detection based on supervised Non-Negative Matrix Factorization, Proc. IEEE Workshop on Biometric Measurements and Systems for Security and Medical Applications (BIOMS), 2011

- [Baluja 2007] S Baluja, H Rowley. Boosting sex identification performance. *International Journal of Computer Vision*, 71:111–119, 2007
- [BenAbdelkader 2002] C BenAbdelkader, R Cutler, L Davis, Person identification using automatic height and stride estimation, *Proc. ICPR*, 2002
- [Bertillon 1889] A Bertillon, Instructions for taking descriptions for the identification of criminals and others, by means of anthropometric indications. American Bertillon Prison Bureau, 1889
- [Bicego 2008] M Bicego, G Brelstaff, L Brodo, E Grosso, A Lagorio, M Tistarelli, Distinctiveness of faces: a computational approach. *ACM Trans. on Applied Perception*, 5(2):1–18, 2008
- [Cao 2008] L Cao, M Dikmen, Y Fu, TS Huang, Gender recognition from body, *Proc. 16th ACM Multimedia*, 2008
- [Cao 2012] D Cao, C Chen, D Adjeroh, A Ross, Predicting gender and weight from human metrology using a copula model, *Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS*, 2012
- [Chen 2012] H Chen, A Gallagher, B Girod, Describing Clothing by Semantic Attributes, *Proc. ECCV*, 2012
- [Chen 2014] C Chen, A Dantcheva, and A Ross, Impact of Facial Cosmetics on Automatic Gender and Age Estimation Algorithms, *Proc. 9th Int. Conf. on Computer Vision Theory and Applications (VISAPP)*, 2014
- [Chen 2015] Q Chen, J Huang, et al, Deep Domain Adaptation for Describing People Based on Fine-Grained Clothing Attributes, *Proc. CVPR 2015*
- [Cole 2007] SA Cole, Twins, Twain, Galton, and Gilman: Fingerprinting, Individualization, Brotherhood, and Race in Pudd'nhead Wilson. *Configurations*, 15(3): 227-265, 2007
- [Collins 2009] M Collins, J Zhang, P Miller, H. Wang, Full body image feature representations for gender profiling, *Proc. ICCV Workshops 2009*, 2009
- [Criminisi 1999] A Criminisi, A Zisserman, L van Gool, New approach to obtain height measurements from video, *Proc. SPIE 3576*, 1999
- [Dangelo 2010] A Dangelo, JL Dugelay, Color based soft biometry for hooligans detection, *Proc. IEEE International Symposium on Circuits and Systems*, 2010
- [Dantcheva 2010] A Dantcheva, JL Dugelay, P Elia, Person recognition using a bag of facial soft biometrics (BoFSB), *Proc. MMSP, IEEE International Workshop on*, pp.511-516, 2010
- [Dantcheva 2011a] A Dantcheva, N Erdogmus, J Dugelay. On the reliability of eye color as a soft biometric trait. *Proc. WACV*, 2011
- [Dantcheva 2011b] A Dantcheva, C Velardo, A D'angelo, JL Dugelay, Bag of soft biometrics for person identification: New trends and challenges, *Multimedia Tools and Applications*, 51(2):739-777, 2011
- [Dantcheva 2015] A. Dantcheva, P. Elia, A. Ross, What else does your biometrics data reveal? A survey on soft biometrics, *IEEE Transactions on Information Forensics And Security (TIFS)*, 2015 (in press)
- [De Angelis 2007] D De Angelis, R Sala, A Cantatore, P Poppa, New method for height estimation of subjects represented in photograms taken from video surveillance systems, *Int. J. Legal Medicine*, 121:489–492, 2007
- [Denman 2009] S Denman, C Fookes, A Bialkowski, Soft-biometrics: Unconstrained authentication in a surveillance environment, *Proc. DICTA*, 2009
- [Denman 2012] S Denman, A Bialkowski, C Fookes, S Sridharan, Identifying Customer Behaviour and Dwell Time using Soft Biometrics, *Video Analytics for Business Intelligence, Studies in Computational Intelligence*, 409: 199-238, 2012
- [Fairhurst 2013] M Fairhurst, *Age Factors in Biometric Processing*, IET Press 2013
- [Farinella 2013] G Farinella, J Dugelay, Demographic classification: Do gender and ethnicity affect each other? *Proc. IEEE/IAPR Int. Conf.. on Informatics, Electronics & Vision, ICIEV 2012*, pp.383,390, 2012
- [Fong 2013] S Fong, Y Zhuang, I Fister, I Fister Jr, A biometric authentication model using hand gesture images, *BioMedical Engineering OnLine*, 12:111, 2013
- [Fu 2010] Y Fu, G Guo, TS Huang, Age Synthesis and Estimation via Faces: A Survey, *IEEE Trans. on PAMI*, 32(11):1955-1976, 2010
- [Fu 2014] S Fu, H He, Z Hou. Learning Race from Face: A Survey. *IEEE Trans. on PAMI*, 36(12):2483-2509, 2014
- [Gallagher 2009] AC Gallagher, T Chen. Understanding images of groups of people. *Proc. Conf. on Computer Vision and Pattern Recognition, CVPR 2009*, pp. 256-263, 2009
- [Ge 2013] Y Ge, J Lu, X Feng, D Yang, Body-based human age estimation at a distance, *Proc. Int. Conf. on Multimedia and Expo Workshops*, 2013
- [Golomb 1990] BA Golomb, DT Lawrence, TJ Sejnowski SEXNET: A Neural Network Identifies Sex From Human Faces. *Proc. NIPS*, 1990
- [Guo 2009a] G. Guo, G. Mu, Y. Fu, and T. Huang. Human age estimation using bio-inspired features. *Proc. Conf. on Computer Vision and Pattern Recognition, CVPR*, pp. 112–119, 2009
- [Guo 2009b] G Guo, G Mu, Y Fu, Gender from body: A biologically-inspired approach with manifold learning. *Proc. of the 9th Asian Conf. on Computer Vision, ACCV 2009*, pp. 236–245, 2009
- [Hadid 2013] A Hadid, M Pietikäinen. Demographic classification from face videos using manifold learning, *Neurocomputing*, 100:197-205, 2013
- [Han 2014] Hu Han, Anil K. Jain, Age, Gender and Race Estimation from Unconstrained Face Images, *MSU Tech. Report: MSU-CSE-14-5*, 2014
- [Hansen 2007] DM Hansen, BK Mortensen, PT Duizer, JR Andersen, TB Moeslund, Automatic Annotation of Humans in Surveillance Video, *Proc. CRV*, pp.473-480, 2007
- [Idrus 2014] SZS Idrus, E Cherrier, C Rosenberger, P Bours, Soft Biometrics Database: A Benchmark For Keystroke Dynamics Biometric Systems, *Proc. BIOSIG*, 2013
- [Itti 1998] L Itti, L, C Koch, W Niebur, A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. on PAMI*, 20(11):1254–1259, 1998
- [Jaha 2014] ES Jaha, MS Nixon. Soft biometrics for subject identification using clothing attributes. *Proc. Int. Joint Conf. on Biometrics, IJCB*, 2014
- [Jaha 2015] ES Jaha, MS Nixon. Viewpoint Invariant Subject Retrieval via Soft Clothing Biometrics. *Proc. Int. Conf. on Biometrics, ICB*, 2015
- [Jain 2004a] AK Jain, SC Dass, K Nandakumar, Soft biometric traits for personal recognition systems, *Proc. Int. Conf. on Biometric Authentication, ICBA*, 2004
- [Jain 2004b] AK Jain, SC Dass, K Nandakumar. Can soft biometric traits assist user recognition? *Proc. SPIE, Biometric Technology for Human Identification*, 5404, 2004
- [Kashyap 2012] AL Kashyap, S Tulyakov, V Govindaraju, Facial behaviour as a soft biometric, *Proc. International Conf. on Biometrics (ICB)*, 2012
- [Kindt 2013] E Kindt, *Privacy and Data Protection Issues of Biometric Applications*, Springer: Law, Governance and Technology Series, 2013
- [Klare 2014] B Klare, S Klum, J Klontz, E Taborsky, T Akgul, AK Jain, Suspect Identification Based on Descriptive Facial Attributes. *Proc. Int. Joint Conf. on Biometrics, IJCB*, 2014
- [Koch 1985] C Koch, and S Ullman, Shifts in selective visual-attention towards the underlying neural circuitry. *Human Neurobiology* 4:219–227, 1985
- [Kumar 2009] Attribute and simile classifiers for face verification N Kumar, AC Berg, PN Belhumeur, and SK Nayar, *Proc. ICCV*, 2009
- [Kwon 1994] YH Kwon, N da Vitoria Lobo. Age classification from facial images. *Proc. IEEE Computer Society Conf. on Computer Vision and Pattern Recognition, CVPR 1994*, pp. 762–767, 1994
- [Lai 2012] K Lai, J Konrad, P Ishwar, Towards Gesture-Based User Authentication, *Proc. Advanced Video and Signal-Based Surveillance (AVSS)*, 2012
- [Lanitis 2002] A Lanitis, C Taylor, T Cootes, Toward automatic simulation of aging effects on face images. *IEEE Trans. on PAMI*, 24(4):442–455, 2002
- [Lanitis 2011] A Lanitis, N Tsapatsoulis, Quantitative evaluation of the effects of aging on biometric templates, *IET Computer Vision*, 5(6) (2011) 338 – 347
- [Lagree 2011] S Lagree, KW Bowyer, Predicting ethnicity and gender from iris texture, *Proc. IEEE Conf. Technologies for Homeland Security*, 2011
- [Layne 2012] R Layne, TM Hospedales, S Gong, Person Re-Identification by Attributes, *Proc. British Machine Vision Conf., BMVC*, 2012
- [Lee 2008] J-E Lee, AK Jain, R Jin, Scars, Marks and Tattoos (SMT): Soft Biometric for Suspect and Victim Identification, *Proc. Biometrics Symposium (BSYM)*, 2008
- [Lei 2013] J Lei, J Zhou, M Abdel-Mottaleb, Gender classification using automatically detected and aligned 3D ear range data, *Proc. Int. Conf. on Biometrics, ICB*, 2013

- [Li 2013] M Li, KJ Han, S Narayanan, Automatic speaker age and gender recognition using acoustic and prosodic level information fusion, *Computer Speech & Language*, 27(1): 151-167, 2013
- [Li 2014] A Li, L Liu, K Wang, S Liu, S Yan, Clothing Attributes Assisted Person Re-identification, in press *Proc. IEEE Transactions on Circuits and Systems for Video Technology*, 25(5): 869-878, 2014
- [Lu 2010] J Lu, YP Tan, Gait-based human age estimation, *IEEE Trans. on IFS*, 5(4): 761 – 770, 2010
- [Lucas 2015] T Lucas, M Henneberg, Comparing the face to the body, which is better for identification? In press *Int. J. Legal Medicine*, 2015, DOI 10.1007/s00414-015-1158-6
- [Lyle 2012] JR Lyle, PE Miller, SJ Pundlik, DL Woodard. Soft biometric classification using local appearance periocular region features, *Pattern Recognition*, 45(11):3877-3885, 2012
- [MacLeod 1994] MD MacLeod, JN Frowley, JW Shepherd. Whole body information: Its relevance to eyewitnesses, *Adult Eyewitness Testimony*, 6, Cambridge University Press, 1994
- [Makihara 2011] Y Makihara, M Okumura, H Iwama, Y Yagi, Gait-based age estimation using a whole-generation gait database, *Proc. Int. Joint Conf. on Biometrics, IJCB*, 2011
- [Makinen 2008] E Makinen, R. Raisamo, An experimental comparison of gender classification methods, *Pattern Recognition Letters*, 29(10):1544–1556, 2008
- [Martinho 2015] D Martinho-Corbishley, MS Nixon, JN Carter, Soft Biometric Recognition from Comparative Crowdsourced Annotations, Accepted *Proc. Int. Conf. on Imaging for Crime Detection and Prevention, ICDP*, 2015
- [Meng 2012] Y Meng, DS Wong, R Schlegel, L Kwok, Touch Gestures Based Biometric Authentication Scheme for Touchscreen Mobile Phones, *Proc. Inscrypt*, 2012
- [Mery 2015] D Mery, KW Bowyer, Automatic facial attribute analysis via adaptive sparse representation of random patches, in press *Pattern Recognition Letters*, doi:10.1016/j.patrec.2015.05.005
- [Merkow 2010] J Merkow, B Jou, M Savvides, An exploration of gender identification using only the periocular region. *Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS* 2010
- [Møgelmoose 2008] A Møgelmoose, TB Moeslund, K Nasrollahi. Multimodal person re-identification using RGB-D sensors and a transient identification database. *Proc. IEEE Workshop Inf., Forensics and Security*, 2013
- [Moghaddam 2002] B Moghaddam, M-H Yang. Learning gender with support faces. *IEEE Trans. on PAMI* 24(5):707–711, 2002
- [Morales 2014] A Morales, J Fierrez, J Ortega-Garcia, Towards predicting good users for biometric recognition based on keystroke dynamics, *Proc. ECCV Workshop on Soft Biometrics*, 2014
- [Mordini 2008] E Mordini, S Massari, Body, biometrics and identity, *Bioethics* 22 (9), 488-498, 2008
- [Ng 2012] CB Ng, YH Tay, B-M Goi. Recognizing Human Gender in Computer Vision: A Survey. 12th Pacific Rim International Conf. on Artificial Intelligence, *PRICAI* 2012, pp. 335-346, 2012
- [Okuma 2004] K Okuma, A Taleghani, N de Freitas, JJ Little, DG Lowe, A Boosted Particle Filter: Multitarget Detection and Tracking, *Proc. ECCV*, 2004
- [Ortega 2009] M Ortega, L Brodo, M Bicego, M Tistarelli. Measuring changes in face appearance through aging. *Proc. Conf. on Computer Vision and Pattern Recognition Workshops, CVPRW*, pp.107-113, 2009
- [Panis 2014] G Panis, A Lanitis. An Overview of Research Activities in Facial Age Estimation Using the FG-NET Aging Database. *Proc. ECCV Workshop on Soft Biometrics*, 2014
- [Park 2009] U Park, A Ross, AK Jain. Periocular biometrics in the visible spectrum: A feasibility study. *Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS*, 2009
- [Park 2010] U Park, AK Jain. Face matching and retrieval using soft biometrics. *IEEE Transactions on IFS*, 5(3):406-415, 2010
- [Park 2011] U Park, RR Jillela, A Ross, AK Jain. Periocular Biometrics in the Visible Spectrum. *IEEE Transactions on IFS*, 6(1):96-106, 2011
- [Ran 2008] Y Ran, G Rosenbush, Q Zheng, Computational approaches for real-time extraction of soft biometrics, *Proc. ICPR*, 2008
- [Rattani 2014] A Rattani, C Chen, A Ross, Evaluation of Texture Descriptors for Automated Gender Estimation from Fingerprints, *Proc. ECCV Workshop on Soft Biometrics*, 2014
- [Reid 2010] DA Reid, MS Nixon, Imputing human descriptions in semantic biometrics, *Proc. ACM Workshop. on Multimedia in Forensics, Security and Intelligence*, 2010
- [Reid 2011] DA Reid, MS Nixon, Using Comparative Human Descriptions for Soft Biometrics, *Proc. Int. Joint Conf. on Biometrics, IJCB*, 2011
- [Reid 2013a] DA Reid, S Samangooei, C Chen, MS Nixon, A Ross, Soft biometrics for surveillance: an overview, *Handbook of Statistics: Machine learning: theory and applications*, 31:327-352, 2013
- [Reid 2013b] DA Reid, MS Nixon, Human identification using facial comparative descriptions, *Proc. Int. Conf. on Biometrics, ICB*, 2013
- [Reid 2014] D Reid, M Nixon, S Stevenage. Soft biometrics; human identification using comparative descriptions. *IEEE Trans. on PAMI*, 36(6):1216-1228, 2014
- [Ricanek 2006] K. Ricanek, T. Tesafaye. MORPH: A longitudinal image database of normal adult age-progression. *Proc. IEEE Int. Conf. on Automatic Face and Gesture recognition, FG2006*, 2006
- [Ruebsamen 2013] T Ruebsamen, J Bayer, C Reich, Mobile Device Biometric Touch Gesture Information Used to Give User Identity Evidence, *Proc. Int. Conf. on Emerging Network Intelligence*, 2013
- [Samangooei 2008] S Samangooei, B Guo, MS Nixon, The use of semantic human description as a soft biometric. *Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS*, 2008
- [Samangooei 2009] S Samangooei, MS Nixon, Performing content-based retrieval of humans using gait biometrics, *Multimedia Tools and Applications* 49(1):195-212, 2009
- [Samangooei 2014] S Samangooei, MS Nixon, On Semantic Soft-Biometric Labels, *Proc. BIOMET 2014*, 2014
- [Sgroi 2013] A Sgroi, KW Bowyer, PJ Flynn, The prediction of old and young subjects from iris texture, *Proc. Int. Conf. on Biometrics, ICB*, 2013
- [Shutler 2002] JD Shutler, MG Grant, MS Nixon, JN Carter, On a large sequence-based human gait database, *Proc. RASC*, 2002
- [Spain 2007] NA Spain, Forensic biometrics from images and video at the Federal Bureau of Investigation. *Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS*, 2007
- [Sultana 2014] M Sultana, PP Paul, M Gavrilova, A Concept of Social Behavioral Biometrics: Motivation, Current Developments, and Future Trends, *Proc. Int. Conf. on Cyberworlds (ICC)*, 2014
- [Takeuchi 2013] A Takeuchi, Y Manabe, K Sugawara, Multimodal soft biometric verification by hand shape and handwriting motion in the air, *Proc. Int. Joint Conf. on Awareness Science and Technology and Ubi-Media Computing (iCAST-UMEDIA)*, 2013
- [Thomas 2007] V Thomas, NV Chawla, KW Bowyer, Learning to predict gender from iris images, *Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS* 2007
- [Tome 2014] P Tome, J Fierrez, R Vera-Rodriguez, M Nixon. Soft Biometrics and their Application in Person Recognition at a Distance. *IEEE Trans. on IFS*, 9(3):464-475, 2014
- [Velardo 2010] C Velardo, J Dugelay, Weight estimation from visual body appearance, *Proc. Int. Conf. on Biometrics: Theory, Applications and Systems, BTAS*, 2010
- [Wallace 2013] L Wallace, JG Trafton, and E Martinson. Complexion as a soft biometric in human-robot interaction, *Proc. BTAS*, 2013
- [Wayman 1997] JL Wayman, Large-scale Civilian Biometric Systems - Issues and Feasibility. *Proc. Card Tech. / Secur. Tech. ID.*, 1997

[Yadav 2013] D Yadav, M Vatsa, R Singh, M Tistarelli. Bacteria Foraging Fusion for Face Recognition across Age Progression. Proc. Conf. on Computer Vision and Pattern Recognition, CVPR, pp. 173-179, 2013

[Yoo 2005] JH Yoo, D Hwang, MS Nixon, Gender classification in human gait using support vector machine, Proc. ACIVS, 2005

[Zhang 2015] H Zhang, JR Beveridge, BA Draper , PJ Phillips, On the effectiveness of soft biometrics for increasing face verification rates, in press CVIU, 2015 <http://dx.doi.org/10.1016/j.cviu.2015.03.003>

Figures and Tables

Fig. 1 Examples of Bertillon’s gathering of measurements [Bertillon 1889].

Fig. 2 Improvement in recognition performance using soft biometrics [Jain 2004a]

Fig. 3 Face attributes and their classification [Kumar 2009]

Fig. 4 Semantic division of facial marks [Park 2010].

Fig. 5 Interface used for labelling semantic descriptions [Samangoeei 2008]

Table 1. Some Semantic Biometric Traits and their Descriptions (Terms)

Body		Global	
Trait	Term	Trait	Term
0. Arm Length	(0.1) Very Short	12. Figure	(12.1) Very Thin
	(0.2) Short		(12.2) Thin
	(0.3) Average		(12.3) Average
	(0.4) Long		(12.4) Big
	(0.5) Very Long		(12.5) Very Big
2. Chest	(2.1) Very Slim	13. Age	(13.1) Infant
	(2.2) Slim		(13.2) Pre Adolescence
	(2.3) Average		(13.3) Adolescence
	(2.4) Large		(13.4) Young Adult
	(2.5) Very Large		(13.5) Adult
3. Figure	(3.1) Very Small		(13.6) Middle Aged
	(3.2) Small		(13.7) Senior
	(3.3) Average	18. Facial Hair	(18.1) None
	(3.4) Large		(18.2) Stubble
	(3.5) Very Large		(18.3) Moustache

Fig. 6 Image highlighting a marked suspect with covered face but wearing distinct clothing, and a face image (bottom right) of a suspect appearing to wear the same clothes [Jaha 2014].

Table 2. Semantic clothing attributes [Jaha 2014]

Body zone	Semantic Attribute	Categorical Labels	Comparative Labels
Head	1. Head clothing category	[None, Hat, Scarf, Mask, Cap]	
	2. Head coverage	[None, Slight, Fair, Most, All]	[Much Less, Less, Same, More, Much more]
	3. Face covered	[Yes, No, Don't know]	[Much Less, Less, Same, More, Much more]
	4. Hat	[Yes, No, Don't know]	
Upper body	5. Upper body clothing category	[Jacket, Jumper, T-shirt, Shirt, Blouse, Sweater, Coat, Other]	
	6. Neckline shape	[Strapless, V-shape, Round, Shirt collar, Don't know]	
	7. Neckline size	[Very Small, Small, Medium, Large, Very Large]	[Much Smaller, Smaller, Same, Larger, Much Larger]
	8. Sleeve length	[Very Short, Short, Medium, Long, Very Long]	[Much Shorter, Shorter, Same, Longer, Much Longer]
Lower body	9. Lower body clothing category	[Trousers, Skirt, Dress]	
	10. Shape	[Straight, Skinny, Wide, Tight, Loose]	
	11. Leg length (of lower clothing)	[Very Short, Short, Medium, Long, Very Long]	[Much Shorter, Shorter, Same, Longer, Much Longer]
	12. Belt presence	[Yes, No, Don't know]	
Foot	13. Shoes category	[Heels, Flip flops, Boot, Trainer, Shoe]	
	14. Heel level	[Flat/low, Medium, High, Very high]	[Much Lower, Lower, Same, Higher, Much higher]
Attached to body	15. Attached object category	[None, Bag, Gun, Object in hand, gloves]	
	16. Bag (size)	[None, Side-bag, Cross-bag, Handbag, Backpack, Satchel]	[Much Smaller, Smaller, Same, Larger, Much Larger]
	17. Gun	[Yes, No, Don't know]	
	18. Object in hand	[Yes, No, Don't know]	
	19. Gloves	[Yes, No, Don't know]	
General style	20. Style category	[Well-dressed, Business, Sporty, Fashionable, Casual, Nerd, Bibes, Hippy, Religious, Gangsta, Tramp, Other]	
Permanent	21. Tattoos	[Yes, No, Don't know]

Fig 7. Division of the body into two different parts for calculating colour histograms to be used as soft biometrics [Møgelmoose 2008].