A New Method for Estimating Body Height from Uncalibrated Images

Anonymous FG2008 submission

Abstract

We address the problem of estimating a person's body height from a single uncalibrated image. The novelty of our work lies in that we handle two difficult cases not previously addressed in the literature: (i) the image contains no reference length in the background scene, (ii) the image contains the upper body only. In a nutshell, our method combines well-known ideas from projective geometry and single-view metrology with prior probabilistic/statistical knowledge of human anthropometry, in a Bayesian-like framework. The method is demonstrated with synthetic (randomly generated) data as well as a dataset of 96 frontal and 12 non-frontal images, and the results are quite encouraging.

1. Introduction

In this paper, we re-visit the problem of body height (or *stature*) measurement from a single *uncalibrated* image. Body height estimation of people in images and video has many important applications, as body height can be used to identify individuals, either uniquely or partially. Furthermore, we focus on applications where the camera calibration parameters are unavailable, such as in forensics and detection/tracking of people from a moving camera or across multiple camera views. In forensic image analysis, body height is used to rule out the possibility that a particular person is the same person in the image (i.e. screening and elimination of suspects) [4, 6, 17, 7]. In human detection/tracking applications, such as video surveillance and customer tracking, body height is used to distinguish among a small set of people in the scene [2, 15, 3, 14, 10].

From a projective geometry point of view, this problem amounts to recovering the length of a 3D line segment based on its projection in a 2D image—the endpoints being the top of the head and the feet. In addition to the projection of the segments endpoints in the image, the solution also requires certain calibration information. In the calibrated case—when full camera calibration parameters are known—we also need to know the direction of the line segment and the plane passing through one of its endpoints [2, 15, 12]. With an uncalibrated image, the necessary calibration information is derived from scene objects depicted in the image and prior knowledge of their 3D structure. One way of achieving this is by computing the length ratio of the line segment with respect to another line segment of known length, called the *reference length*. In the most general case of (strong) perspective projection, this approach is applicable under certain circumstances only, however, such as the following: (i) the two line segments are coplanar and the vanishing line of their plane is known; or (ii) the two line segments are parallel and the vanishing points of their direction and that of the line passing through an endpoint of each line segment are known; or (iii) the two line segments are collinear and the vanishing line of their direction is known [17, 9, 8, 13]. The reference length and vanishing points/lines are obtained from prior knowledge of the 3D structure of the background scene. Note that in the case of weak perspective (i.e. if the line segments are coplanar and their plane is almost parallel to the image plane), the length ratio of the line segments in the image is equal to their length ratio in 3D, and hence only the reference length is needed. So, whichever the camera model assumed, a reference length is always required by this geometric approach for computing the length of a 3D line segment from an uncalibrated image.

Recently, we have developed a novel visual metrology technique that extends the latter described approach in one important way: a reference length is no longer required. In a nutshell, this is achieved by incorporating certain probabilistic/statistical properties of human anthropometry into the estimation process [1]. In the present work, we apply this technique to body height measurement, and we handle two difficult cases not previously addressed in the literature as far as body height estimation: (i) the case where the image contains no reference length in the background scene, and (ii) the case where the image contains the upper body part only (rather than the whole body).

2. Methods

We start by presenting the visual metrology technique that we developed recently [1] (Section 2.1), then we discuss how this technique is applied to body height estimation from full-body images (Section 2.2) and upper-body images

of people (Section 2.3). Finally, we discuss some details about the implementation of these methods in Section 2.4.

2.1. Probabilistic Visual Metrology

Let $\mathbf{x}$ be a vector containing the lengths of n 3D line segments, and let $\mathbf{y}$ be a vector containing m length ratios x_i/x_j where $i \neq j$. The goal is to estimate $\mathbf{x}$ based on the observed/known value of $\mathbf{y}$.

From a purely algebraic point of view, this is a **ill-posed inverse problem**, as there are infinitely many vectors $\mathbf{x}$ that could generate the same value of $\mathbf{y}$ (at the very least, if a solution $\mathbf{x}_0$ exists, then *any* scalar multiple of $\mathbf{x}_0$ is also a solution). In the sequel, we will show that it is possible to solve this inverse problem by introducing appropriate prior statistical knowledge about $\mathbf{x}$ in a Bayesian-like framework.

Let $\mathcal{I}$ be the set of pairs (i,j) such that the ratio x_i/x_j is included in $\mathbf{y}$. Let r_{ij} be the observed value of x_i/x_j, for any $(i,j) \in \mathcal{I}$. Each r_{ij} value provides a linear equality constraint in the unknowns x_i and x_j:

$$x_i - r_{ij}x_j = 0 \tag{1}$$

Hence with m distinct length ratios we obtain m corresponding linear equality constraints, which we express in matrix form as a linear homogeneous system:

$$\mathbf{C}\mathbf{x} = \mathbf{0} \tag{2}$$

Barring numerical errors, r_{ij} and r_{ji} provide equivalent equality constraints, and $\mathbf{C}$ hence has at most $n(n-1)/2$ rows. We know from basic linear algebra that an exact non-trivial solution for the system in (2) exists if and only if $\mathbf{C}$ is rank-deficient, i.e. $\mathrm{rank}(\mathbf{C}) < n$. Otherwise, when $\mathrm{rank}(\mathbf{C}) = n$, a non-exact nontrivial solution can be obtained by solving the following minimization problem via least squares estimation:

$$\min_{\mathbf{x}} \|\mathbf{C}\mathbf{x}\|^2 \qquad \text{subject to} \quad \|\mathbf{x}\|^2 = 1 \tag{3}$$

The above LSE problem has a unique solution, which happens to be the right singular vector, denoted $\mathbf{v}_n$, of $\mathbf{C}$ that corresponds to its smallest singular value [18]. Thus, a set of approximate (non-exact) solutions for (2) is:

$$\mathbf{x} = k\mathbf{v_n} \tag{4}$$

where k is an arbitrary scalar. But clearly this is an infinite and unbounded solution space, which confirms our earlier conjecture that the inverse problem in (2) is ill-posed. Some sensible way is thus needed of selecting a single "best" solution from this infinitum of candidate solutions.

A well-known strategy for dealing with this situation is *regularization*, which generally involves imposing additional constraints derived from prior knowledge about the unknown $\mathbf{x}$ [19]. For this, we shall assume $\mathbf{x}$ is a random vector (equivalently each x_i is a random variable), and use the following statistical properties to formulate constraints: (i) lower and upper bounds of each component x_i of $\mathbf{x}$, (ii) lower and upper bounds of each component ratio x_i/x_j, and (iii) the prior probability distribution of $\mathbf{x}$.

A few definitions first. Let $L_{i,i}$ and $U_{i,i}$ be respectively the lower and upper α percentiles[1] of the probability distribution of x_i. Similarly, let $L_{i,j}$ and $U_{i,j}$ respectively be the lower and upper bounds of the quotient random variable $\rho_{ij} = x_i/x_j$, for any $i \neq j$. Clearly, with a sufficiently small α, α percentiles act as de facto lower and upper bounds of the respective random variable. We also assume that the prior probability of $\mathbf{x}$, denoted $\Pi(\mathbf{x})$, has a Gaussian distribution $N(\mu, \mathbf{\Sigma})$, and we define the log prior function as:

$$\ell(\mathbf{x}) = (\mathbf{x}-\mu)^T\mathbf{\Sigma}^{-1}(\mathbf{x}-\mu) \propto \log \Pi(\mathbf{x}) \tag{5}$$

Now, if the values of $L_{i,j}$ and $U_{i,j}$ are known for all $(i,j) \in \mathcal{I}$, then the following set of inequality constraints on $\mathbf{x}$ are obtained:

$$L_{i,i} \leq x_i \leq U_{i,i} \qquad \text{for all } i \in \{1, 2, \cdots, n\} \tag{6}$$

$$L_{i,j} \leq \frac{x_i}{x_j} \leq U_{i,j} \qquad \text{for all } (i,j) \in \mathcal{I} \tag{7}$$

By imposing these constraints on our inverse problem, we filter out unlikely solutions and narrow down the infinite and unbounded solution space in (4). As Figure 1 illustrates, this amounts to a bounded (but still infinite) solution space. Furthermore, based on the prior probability distribution of $\mathbf{x}$, a unique solution to our inverse problem can be obtained as the most probable point in this new solution space. More formally, we re-formulate (3) as follows:

$$\min_{\mathbf{x}}[\|\mathbf{C}\mathbf{x}\|^2 + \lambda\ell(\mathbf{x})] \ \text{ subject to } \begin{cases} L_{i,i} \leq x_i \leq U_{i,i} \\ L_{i,j} \leq \frac{x_i}{x_j} \leq U_{i,j} \end{cases} \tag{8}$$

where λ is a regularization parameter. The cost functional is now a weighted sum of two terms: (i) a *data term* that has the effect of enforcing observed data, and (ii) a *prior term* that has the effect of favoring more probable candidate solutions. Since the inequalities in (7) can be expressed as linear inequalities in x_i and x_j, this is a **linearly** constrained quadratic function minimization problem—also called *quadratic programming* (QP). We use standard iterative QP techniques to solve this quadratic program (currently the `quadprog` Matlab function). In [1] we also discuss a more efficient non-iterative variant of this technique. Interestingly, assuming input errors are of the form:

[1] The lower α percentile of the probability distribution of a random variable X is a value z_L such that $\Pr(X \leq z_L) = \alpha$. Similarly, its upper α percentile is a value z_R such that $\Pr(X \leq z_R) = 1 - \alpha$.

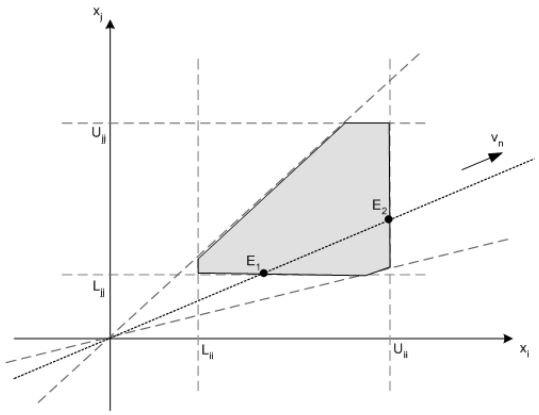

Figure 1. Geometric interpretation of the solution space of our inverse problem when $n = 2$. Initially (before regularization), the solution space consists of an unbounded one-dimensional subspace (*dotted line passing through origin*). The range constraints (*dashed lines*) define a sort of feasibility region (*filled polygon*). By further applying the range constraints, we obtain a bounded one-dimensional subspace (*line segment between* E_1 and E_2).

$\mathbf{Cx} = \eta$, where η is Gaussian white noise, the solution of (8) in fact corresponds to the Bayesian MAP (maximum a posteriori) estimate of the inverse problem in (2).

2.2. Body Height from a Whole-Body Image

Given a whole-body photograph of a person, with the person in upright standing pose (not slouching or leaning), the metrology technique of Section 2.1 can be used to estimate his/her body height as follows; the n line segments consist of a set of anthropometrics including body height, and each anthropometric is the line segment between two visible and well-defined body landmarks. Furthermore, if the image is uncalibrated, we need to choose these anthropometrics such that their ratios can be easily computed from an image, which means: (i) they are either collinear or coplanar, and (ii) the required vanishing points and lines can be inferred from the image (see Section 1).

Assuming the person in the image is in frontal pose with the head is in neutral pose, we use the following set of **ten anthropometrics**[2]: (1) body height, (2) neck height, (3) acromial height, (4) head to chin distance, also called *head length*, (5) stomion to top of head distance, (6) subnasale to top of head distance, (7) forehead to chin distance, (8) sellion to chin distance, (9) biocular distance, i.e. between outer corners of the eyes, (10) bitragion distance. These anthropometrics are illustrated in Fig. 2. In addition to being all vertical, we also define anthropomerics 1–8 to be collinear along the midline of symmetry (see Section 2.4), so that we need at most the vertical vanishing point to compute their pairwise ratios. Anthropometrics 4–10 lie on the facial surface and are nearly coplanar, assuming the person is sufficiently far from the camera. Furthermore, since the person is assumed to be in frontal pose, the weak perspective model is a good approximation over the facial region and we do not need any calibration information to compute the pairwise ratios of these anthropometrics. Currently we compute pairwise ratios of anthropomerics 1–8 and those of anthropomerics 4–10, hence a total of (8.7)/2+(7.6)/2=49 ratios (so $n = 10$ and $m = 49$).

When the person is in a non-frontal pose (with the whole body and/or the head at an angle with respect to the camera), however, the facial anthropometrics no longer lie in a plane parallel to the image plane, and so to compute their pairwise ratios we would need a lot more information than what is available in an uncalibrated image. Consequently we only use anthropometrics 1–8 for such images.

Clearly there are many other possible anthropometrics that can be used in this body height estimation method, in addition to or in place of those we chose in this paper. We do not claim that our choice is by any means optimal (in terms of estimation accuracy). It is not arbitrary, however, and is motivated by three key factors: (i) ability to compute the anthropometric ratios from an uncalibrated image, (ii) the anthropometric is correlated with body height, (iii) availability of the required anthropometric statistics, and (iv) accuracy of landmark localization in the image. For example, the *bigonial width*[3] is a bad choice because the gonial landmarks are difficult to locate in an image, whether manually or automatically, and even with high-resolution images. The same can be said about *arm length*, since a person's arm is often times bent and/or occluded by clothing.

2.3. Body Height from an Upper-Body Image

We extend the method of the previous subsection to handle images containing only the upper body part of a person. Specifically, we first estimate some set of upper body anthropometrics using the same technique, and then we compute body height as a linear function of the estimated anthropometrics. The parameters of the linear prediction model are learned from a large set of randomly generated data (Section 3.1). The main challenge with this approach is that apart from facial anthropometrics, there are pretty much no anthropometrics that are at once: easily obtainable from images, are strongly correlated with body height, and whose statistics are available. We use the acromion to top of head distance, in addition to the seven facial anthropometrics from the previous subsection. Because anthropometric statistics for the former are not available in any well-known

[2] *trapezius*: the point at the side of the neck, *acromion*: the tip of the shoulder, *subnasale*: the point just below the nose, *tragion*: the point on the cartilaginous flaps in front of each earhole, *stomion*: the point at the center of the mouth, *sellion*: at the deepest point of the nasal root depression (also called nose bridge) [11, 16].

[3] Bigonial (or mandible) width: the straight-line distance between the left and right *gonial* landmarks—the corners of the jaw.

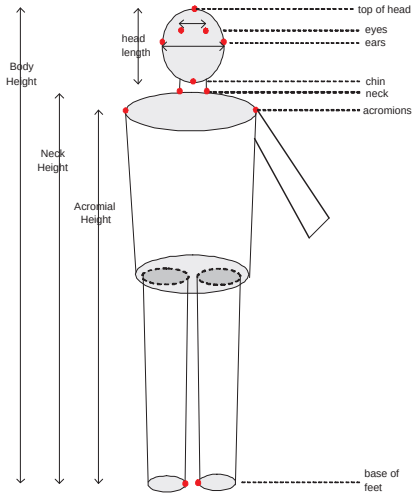

Figure 2. Anthropometrics we estimate in this paper.

anthropometric surveys, we have instead derived them from the statistics of body height and acromial height.

2.4. Method Implementation Details

Localization of the body landmarks associated with these anthropometrics is achieved as follows (Fig. 3). For frontal-view images, we start by locating the following 13 body landmarks in the image: top of the head, forehead, subnasale, stomion, chin, left and right corners of the eyes, left and right tragions, left and right trapezius, left and right acromions, and left and right medial longitudinal foot arches. This is currently done semi-automatically by having the user select points in the image via an interactive Matlab interface. We then estimate the person's medial axis (midline of symmetry) as the line passing through the top of the head, midpoint of the two foot landmarks, and the vertical vanishing point. We refine the locations of the first five landmarks (top of head, forehead, subnasale, stomion, chin) by projecting them onto the medial axis. For the left/right landmark pairs of the trapezius, neck, acromions, and foot arches, we compute the intersection of the medial axis with the line segment joining each landmark pair. This way, the vertical anthropometrics are all collinear with the medial axis. Landmark localization in side-view images is the same way, except that we do not locate the eye corners and tragions, and we estimate the midline of symmetry as the line passing through the top of the head, midpoint of the two foot landmarks, and the vertical vanishing point.

The anthropometric statistics required for this method (namely the lower and upper bounds of both each anthropometric and each anthropometric ratio, and the mean and covariance of all anthropometrics) are obtained from the seminal anthropometric survey in [5]. In particular, the upper and lower bounds of each anthropometric are obtained as

(a) (b)

Figure 3. Landmark localization in (a) frontal-view images, and (b) side-view images. The *blue dots* are points we mark manually; *green dotted line* is the medial axis of symmetry of the person, which we estimate by fitting a line to some of those dots (see text)

the upper and lower 100α percentiles of the normal distribution $N(\mu, \sigma)$, where μ and σ are respectively its mean and standard deviation. We should also note that these statistics are **categorized by gender**, i.e. separate statistics for males and females [1].

3. Experiments and Results

3.1. The Data

We tested our body height estimation methods both on synthetic data and real images. The former consists of **2000** random vectors $\mathbf{y}$, each of which is obtained by first randomly generating a vector $\mathbf{x}$ from the multivariate Gaussian distribution of male or female anthropometrics, then computing the corresponding pairwise ratios (Fig. 4 (a)). The parameters of the anthropometric distributions are obtained from the anthropometric survey in [5]. Notice that the generated $\mathbf{x}$ and $\mathbf{y}$ vectors are error-free.

The real images consist of a set of high-resolution (4368x2912) images captured in-house using a Canon 28-200mm EOS camera (Fig. 4 (b)). Each image is a full-body shot of one person. The dataset contains a total of 108, of 27 different adults, 7 females and 20 males, spanning various ethnicities: caucasian, chinese, indian, and african. Also, 96 of the images are frontal-view shots and 12 images are side-view shots, all taken at about a 45-degree angle. The vertical vanishing point needed to compute pairwise ratios of anthropometrics 1–8 is computed as the intersection in the image of parallel vertical lines of the background scene.

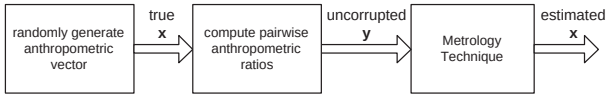

(a)

(b)

Figure 4. (a) Methodology for (randomly) generating synthetic data, (b) Sample images from our inhouse dataset.

3.2. Results

As discussed, we estimate body height by applying our novel metrology technique (Section 2.1) to estimate a set of anthropometrics. For whole-body images, this set includes body height (Section 2.2), while for upper-body images it does not, and a linear model is subsequently applied to predict body height (Section 2.3). To simplify the discussion, let us assign a number to each anthropometric, as follows: (1) body height, (2) neck height, (3) acromial height, (4) head to chin distance, (5) stomion to top of head distance, (6) subnasale to top of head distance, (7) forehead to chin distance, (8) sellion to chin distance, (9) biocular distance, i.e. between outer corners of the eyes, (10) bitragion distance, (11) acromion to top of head distance.

We compute body height estimation error as: `estimated value - true value`. For the method parameters, we use $\alpha = 1e-7$ and $\lambda = 0.01$. Just to have an idea, the lower and upper bounds for body height with this α value are respectively 1.29m and 2.22m for males, and 1.18m and 2.07m for females.

In order to investigate the effect of choice of the anthropometric set on estimation error, we tested the method with various different sets. The two plots in Fig. 5 show the mean, minimum, and maximum of estimation error for each of the following anthropometric sets:

1. $\{1, 2\}$,
2. $\{1, 3\}$,
3. $\{1, 2, 3\}$,
4. $\{1, 11\}$,
5. $\{1, 4, 5, 6, 7, 8\}$,
6. $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$,
7. $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$,
8. $\{4, 5, 6, 7, 8\}$,
9. $\{9, 10\}$,
10. $\{4, 5, 6, 7, 8, 9, 10\}$,
11. $\{4, 5, 6, 7, 8, 9, 10, 11\}$.

Sets 1–7 correspond to body height estimation from full-body images. The best performance is obtained with sets #6 and #7, which suggests that the acromion to top of head distance does not add much new information to the estimation process. Sets 8–11 are used to *simulate* body height estimation from upper-body images (since all our images actually are whole-body images). The best performance is obtained with set #11, albeit interestingly the performance of the other three sets (#8, #9 and #10) is not much different. Two other important observations are in order: (1) estimation error with sets 1–7 is significantly better than with sets 8–11, which makes sense, and (2) estimation error is significantly larger for real images than synthetic data, which is not surprising since the latter contains no input error. In fact, the estimation error obtained with synthetic data is equal to the *model error* of the method.

Fig. 6 shows estimation results for anthropometric sets #6 and #11 in more detail, based on real frontal-view images in both cases. The mean, standard deviation, minimum, and maximum values are respectively (in cm): $-\mathbf{5.3}, \mathbf{6.3}, -\mathbf{20.1}, \mathbf{11.7}$ for set#6, and $-\mathbf{4.0}, \mathbf{8.3}, -\mathbf{28.6}, \mathbf{13.0}$ for set#11. Plot (b) suggests a mild negative correlation between estimation error and the true value, and in fact, the computed correlation coefficient is -.50 for set#6 and -.82 for set#11. This implies that estimation error tends to be larger for people with "extreme" body height, i.e. are far from the population mean (either too tall or too short). This should not be surprising since our metrology technique favors candidate solutions with larger prior probability, hence close to the population mean.

In order to investigate the effect of camera viewpoint (i.e. camera angle with respect to person) on estimation performance, we tested the method with 12 side-view images and 12 frontal-view images of the same set of 6 people (i.e. 2 images per person per view). The side-view images are all at about a 45-degree angle. The results are shown in Fig. 7. The mean, standard deviation, minimum, and maximum values are respectively (in cm): $-\mathbf{3.4}, \mathbf{11.0}, -\mathbf{23.1}, \mathbf{14.1}$ for frontal-view images, and $-\mathbf{6.1}, \mathbf{11.8}, -\mathbf{28.1}, \mathbf{11.1}$ for side-view images. Interestingly, the standard deviation and range of estimation error (i.e. difference between minimum and maximum) is the almost same in both cases. Also, as

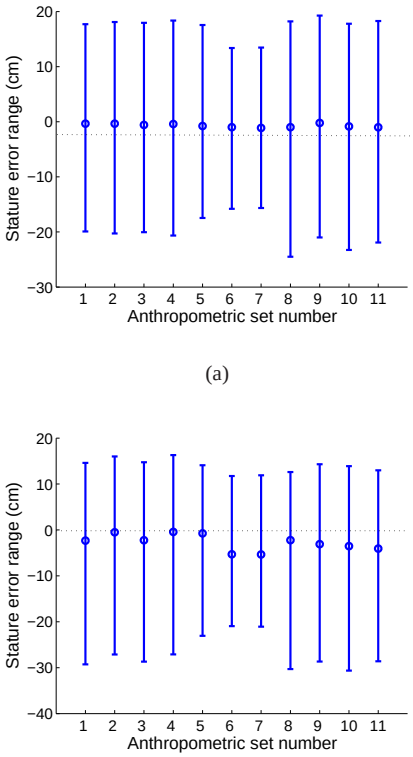

Figure 5. Mean and range of body height estimation error based on different anthropometric sets (see text): (a) with synthetic data, (b) with 96 real frontal-view images.

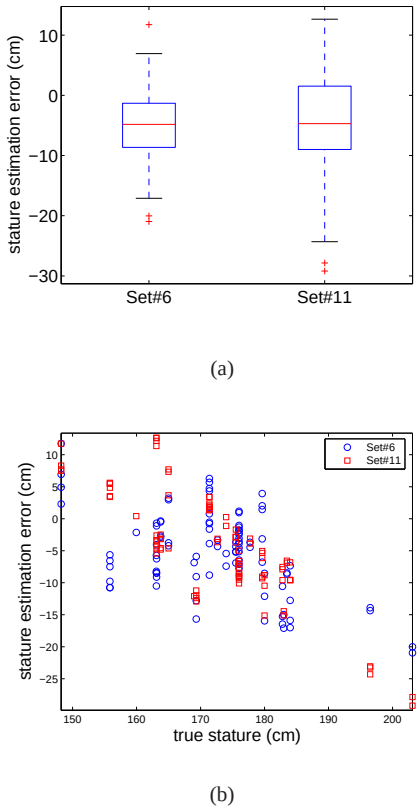

Figure 6. Body height estimation results for anthropometric sets#6 and #11 and based on real frontal-view images: (a) boxplot of estimation error, (b) estimation error vs. true body height.

with Figure 6, plot (b) here indicates a mild negative correlation between estimation error and the true value for both frontal and side-view images. The computed correlation coefficient is -.89 and -.96 respectively for frontal-view and side-view images. Given the small sample size, these results are at best inconclusive. Nonetheless, they seem to suggest that there is no significant difference between estimation from frontal-view and side-view images except for an overall shift in estimation error of about 7cm. Further investigation with a larger dataset is needed to confirm these preliminary observations.

4. Conclusions and Future Work

We presented methods for estimating a person's body height from a single uncalibrated image. The novelty is two-fold: (i) we do not require scale information (the so-called *reference length*); and (ii) we handle images containing the upper-body part only of the person. The former is achieved via clever use of prior knowledge of anthropometric statistics. The latter is achieved by extending whole-body height estimation with a linear prediction model.

Based on tests with synthetic data and real images, performance was found to be inferior to that of other well-known single-view metrology techniques such as [9]. But of course this is the price to be paid for substituting the reference length with statistical information, and for not using whole-body images. Furthermore, inaccurate estimates are not useless; they can be combined with other (inaccurate) sources of information to yield reliable estimates. Another limitation of our method is that it is currently semi-automatic. Full automation would mean automating both gender classification and landmark localization in the image. The latter is no big deal, while the former is an open problem in computer vision.

In future work, we plan to: (i) extend method to estimation from video (a sequence of image); at the very least, this will help reduce estimation inaccuracy caused by landmark localization error; (ii) investigate the effect of camera viewpoint (i.e. camera angle with respect to person) on estimation accuracy in a more comprehensive and systematic manner; (iii) identify the major sources of error in our body height estimation methods; (iv) derive analytical and empirical measures that characterize the sensitivity of these methods to random input error; (v) explore other prediction models, particularly non-linear ones, for estimating body

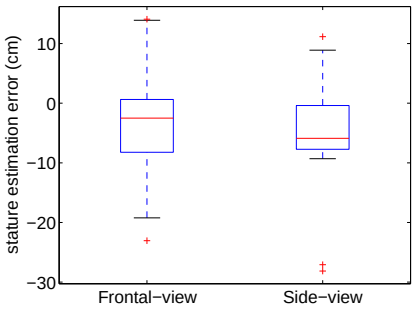

(a)

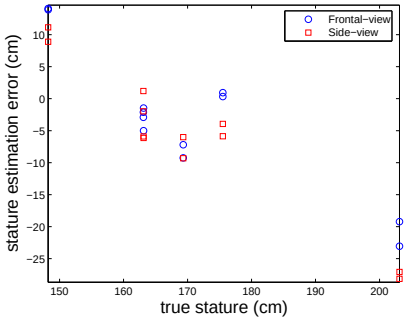

(b)

Figure 7. Body height estimation results based on 12 frontal-view and 12 side-view images: (a) boxplot of estimation error, (b) estimation error vs. true body height.

height from upper-body images.

References

[1] Anonymous submission to ECCV 2008. 1, 2, 4

[2] C. BenAbdelkader, R. G. Cutler, and L. S. Davis. View-invariant estimation of height and stride for gait recognition. In *Workshop on Biometric Authentication*, 2002. 1

[3] A. Bovyrin and K. Rodyushkin. Human height prediction and roads estimation. In *IEEE Conference on Advanced Video and Signal-Based Surveillance*, 2005. 1

[4] S. Bramble, D. Compton, and L. Klasen. Forensic image analysis. In *13th Interpol Forensic Science Symposium*, 2001. 1

[5] J. Cheverud, C. G. Gordon, R. Walker, C. Jacquish, L. Kohn, A. Moore, and N. Yamashita. 1988 anthropometric survey of us army personnel. Technical Report TR-90/032, US Army Natick Research Labs, May 1990. 4

[6] D. Compton, C. Prance, M. Shears, and C. Champod. A systematic approach to height interpretation from images. In *Proc. of SPIE*, volume 4232, 2001. 1

[7] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. In *CVPR*, pages 434–442, Sept. 1999. 1

[8] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. *IJCV*, 40(2):123–148, Nov. 2000. 1

[9] A. Criminisi, A. Zisserman, L. Van Gool, S. Bramble, and D. Compton. A new approach to obtain height measurements from video. In *Proc. of SPIE*, volume 3576, 1998. 1, 6

[10] D. DeAngelis, R. Sala, A. Cantatore, P. Poppa, M. Dufour, M. Grandi, and C. Cattaneo. New method for height estimation of subjects represented in photograms taken from video surveillance systems. *International Journal of Legal Medicine*, 121(6):489–492, November 2007. 1

[11] L. G. Farkas, editor. *Anthropometry of the Head and Face*. Raven Press, second edition, 1994. 3

[12] D. Hansen, B. Mortensen, P. Duizer, J. Andersen, and T. Moeslund. Automatic annotation of humans in surveillance video. In *Fourth Canadian Conference on Computer and Robot Vision*, pages 473–480, 2007. 1

[13] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition, 2004. 1

[14] A. Leykin and M. . Tuceryan. A vision system for automated customer tracking for marketing analysis: low-level feature extraction. Technical Report TR612, Indiana University, June 2005. 1

[15] C. Madden and M. Piccardi. Height measurement as a session-based biometric. In *Image and Vision Computing New Zealand*, 2005. 1

[16] S. Pheasant. *Bodyspace: Anthropometry, Ergonomics and the Design of Work*. Taylor and Francis, second edition, 1996. 3

[17] N. Saitoh, K. Kurosawa, and K. Kuroki. A study on height measurement from a single view. In *ICIP*, 1999. 1

[18] G. Strang. *Introduction to Linear Algebra*. Wellesley-Cambridge Press, third edition, 2003. 2

[19] C. R. Vogel. *Computational Methods for Inverse Problems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2002. 2