

Statistical Body Height Estimation from a Single Image

Chiraz BenAbdelkader
New York Institute of Technology
Abu Dhabi, United Arab Emirates
chiraz@nyit.edu

Yaser Yacoob
Institute of Advanced Computer Studies
University of Maryland, College Park, MD USA
yaser@umiacs.umd.edu

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 to indicate absolute scale, (ii) the image contains the upper body part 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 images.

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. Applications where the camera calibration parameters are unavailable include forensics and detection/tracking of people from a moving camera or across multiple camera views. In forensic image analysis, body height can be 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, 18, 7]. In human detection/tracking applications, such as video surveillance and customer tracking, body height can be used to distinguish among a small set of tracked people in the scene [1, 16, 3, 15, 10].

From a projective geometry point of view, this amounts to the basic 3D reconstruction problem of recovering the length of a 3D line segment based on its projection in a 2D image. In the calibrated case, this problem can be solved with additional knowledge of the segment's direction and a plane or line passing through one of its endpoints [1, 16, 13]. A typical scenario consists of the line segment being orthogonal to the "ground plane", and one of the end-

points lies on this plane, such as in [1]. In the uncalibrated case, a general approach for this problem consists of first computing the length ratio q of the line segment with respect to another line segment in the scene of known length L , called the *reference length*, and then simply obtaining the solution as the product ($q \cdot L$). This approach is purely geometric and comprises a whole family of metrology techniques [18, 9, 8, 14, 12]. However, it poses two main difficulties in practice: (i) a reference length may not always be available, and (ii) under the general (strong) perspective camera model, the length ratio q can only be computed under certain circumstances, such as when the two line segments are coplanar and the vanishing line of their plane is known, or when the two line segments are collinear and the vanishing point of their direction is known.

The uncalibrated scenario is what we are concerned with in this paper. Recently, we have developed a novel visual metrology technique that extends the above described approach for this scenario in one important way: it obviates the need for a reference length [2]. This is achieved by incorporating certain probabilistic/statistical properties of human anthropometry into the estimation process. In the present work, we apply this same technique to body height measurement from whole body and upper body images of people. In other words, the input to our estimation methods consists of anthropometric ratios (which are measured from the image) and anthropometric statistics. Our main contribution lies in that we handle two difficult cases not previously addressed in the literature as far as body height estimation from uncalibrated images: (i) the image contains *no* information on absolute scale (namely a reference length), and (ii) the image contains the upper body part only.

The main premise of our body height estimation method can be explained in an information-theoretic sense as follows. Basically, since human body dimensions and proportions (i) vary within a relatively limited range and (ii) are correlated, the projection of a person in an image must encode *some* information about the absolute scale of the imaged scene. Therefore, we expect to be able to estimate the person's body dimensions, such as body height, to within

accuracy better than a random guess but less than when accurate/complete absolute scale information is available.

2. Methods

We start by presenting the visual metrology technique that we developed recently [2] (Section 2.1), then we discuss how this technique can be 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. Statistical Visual Metrology

We describe a novel metrology technique that simultaneously estimates *multiple*, say n , 3D line segments (such as anthropometrics) from a single image under minimal calibration information, with $n > 1$. 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 an **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 find an approximate solution 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 :

$$x_i - r_{ij}x_j = 0 \quad (1)$$

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

$$\mathbf{C}\mathbf{x} = \mathbf{0} \quad (2)$$

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

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

The condition $\|\mathbf{x}\|^2 = 1$ serves both to avoid the trivial solution and to guarantee uniqueness. This 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 [19]. Thus, a set of approximate (non-exact) solutions for (2) is:

$$\mathbf{x} = k\mathbf{v}_n \quad (4)$$

where k is an arbitrary scalar. This is an infinite and unbounded solution space, with infinite candidate solutions. Furthermore, since $\|\mathbf{C}\mathbf{x}\|^2 = k^2\|\mathbf{C}\mathbf{v}_n\|^2 = k^2\sigma_n^2$, the data fit (namely the sum of square errors) of the solutions is an increasing function of k , and the trivial solution has the best fit. Thus, data fit by itself does not seem to be a sensible way for selecting a single "best" solution from among these candidate solutions.

As mentioned earlier, this situation arises because fundamentally this is an underdetermined inverse problem; there exist infinite \mathbf{x} 's that can generate the same input \mathbf{y} . Hence what we need to do is impose additional constraints on \mathbf{x} in order to narrow down the solution space. This is generally called *regularization*. To this end, we shall assume \mathbf{x} is a random vector (equivalently each x_i is a random variable), and formulate constraints based on the following statistical properties of \mathbf{x} : (i) range of variation of both \mathbf{x} and \mathbf{y} , and (ii) 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 percentiles¹ 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 , the 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 $N(\mu, \Sigma)$, and we define the log prior function (also called *mahalanobis distance*) as:

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

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:

$$L_{i,i} \leq x_i \leq U_{i,i} \quad \text{for all } i \in \{1, 2, \dots, n\} \quad (6)$$

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

Furthermore, we can include a smoothing term in the cost functional of the minimization problem in (3) that favors solutions with higher prior probability, i.e. smaller $\ell(\mathbf{x})$.

¹Let X be a random variable with mean μ and standard deviation σ . The lower α percentile of X is the value x_L such that $\Pr(X \leq x_L) = \alpha$, and its upper α percentile is the value x_R such that $\Pr(X \leq x_R) = 1 - \alpha$. Furthermore, if X has a normal distribution, then $x_L = \mu - t \cdot \sigma$ and $x_R = \mu + t \cdot \sigma$, where $t = \chi^{-1}(1 - 2\alpha)$.

By doing so, in effect we incorporate knowledge about the correlatedness of the \mathbf{x} variables (in addition to knowledge about their range of variation). Intuitively, this knowledge includes things like a person taller than 1.80m is unlikely to have a head length of less than 12cm. We thus re-formulate our inverse problem of (2) as follows:

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

where λ is a regularization parameter that controls the relative effect of the data term and smoothing term. It is easy to see that this is in effect a **linearly** constrained quadratic function minimization problem (also called *quadratic programming* (QP)), since the inequalities in (7) can be expressed as linear inequalities in x_i and x_j . We use standard iterative QP techniques to solve this quadratic program (currently the `quadprog` Matlab function).

Interestingly, assuming input errors are of the form: $\mathbf{C}\mathbf{x} = \eta$, where η is Gaussian white noise, then our solution of (3) corresponds to the Maximum Likelihood estimate (MLE), and the solution of (8) corresponds to the Bayesian maximum a posteriori (MAP) 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), we will show that the metrology technique of Section 2.1 can be used to estimate his/her body height. Namely, the n line segments consist of a set of anthropometrics including body height, and each anthropometric is the straight line distance between two visible and well-defined body landmarks (canonical points). Furthermore, because the image is uncalibrated, these anthropometrics are chosen such that their ratios can be measured from the image, namely: they need to be either collinear or coplanar, and the vanishing points/lines, if needed, can be readily measured from the image (see Section 1).

The premise of this body height estimation method can be formulated in an information-theoretic sense as follows. Since human body dimensions and proportions (i) vary within a relatively limited range and (ii) are correlated, the projection of a person in an image must encode *some* information about the absolute scale of the imaged scene. Therefore, we expect to be able to estimate the person’s body dimensions, such as body height, to within accuracy better than a random guess but less than when accurate/complete absolute scale information is available.

We currently use the following set of **ten anthropometrics**²: (1) body height, (2) trapezium (or neck) height, (3)

²*trapezium*: the point at the side of the neck, *acromion*: the tip of the

acromial height, (4) head to chin distance (or 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. 1.

For the first four anthropometrics the person is assumed to be in upright standing pose—not slouching or leaning on one side. Anthropometrics 1–8 are all vertical and collinear along the midline of symmetry, and so we only need the vertical vanishing point to compute their pairwise ratios. Anthropometrics 4–10 lie on the facial surface and are nearly coplanar, provided that the person is sufficiently far from the camera. Furthermore, when the head pose is nearly frontal to the camera, the weak perspective model is a good approximation over the facial region, and in this case computation of the pairwise ratios of facial anthropometrics requires no calibration information. Currently we compute pairwise ratios of anthropometrics 1–8 and pairwise ratios of anthropometrics 4–10, hence a total of 39 ratios.

However, if the whole body and/or the facial plane are at an angle with respect to the camera, then the weak perspective assumption no longer applies, and so computing pairwise ratios of the facial anthropometrics (4–10) requires a lot more information than what is available in an uncalibrated image. Consequently we only use anthropometrics 1–8 for such images.

Finally, a word on the choice of anthropometrics. In theory, there are many other possible anthropometrics that can be used, in addition to or in place of the ones we use in this paper. However, some factors have limited this choice, in particular: (i) ability to compute the anthropometric ratios from an uncalibrated image, (ii) availability of anthropometric statistics, and (iii) accuracy of landmark localization in the image; for example, the bigonial width³ is a bad choice because the gonial is difficult to locate in an image, even manually and even in 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 Section 2.2 to handle images containing only the upper body part of the person. In a nutshell, we first estimate some set of upper body anthropometrics using the same technique, and then we use a linear prediction model to compute (or predict) body height from the estimated anthropometric values. The model parameters are learned from a large set of randomly generated data

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, 17].

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




Figure 1. Anthropometrics we estimate in this paper.

(Section 3.1). The main challenge we face with this approach lies in that, aside from facial anthropometrics, there are pretty much no anthropometrics that are at once: easily obtainable from images, strongly correlated with body height, and whose statistics are available (see discussion at the end of Section 2.2). Currently we use the seven facial anthropometrics from Section 2.2 (4–10), plus the acromion to top of head distance. Because statistics for the latter anthropometric are not available in any well-known anthropometric surveys, we have instead derived them from the statistics of body height and acromial height.

2.4. Implementation Details

Localization of the body landmarks associated with the anthropometrics of concern is achieved as follows (Fig. 2). For frontal-view images, we start by locating these 15 body landmarks in the image: top of the head, forehead, subnasale, stomion, chin, left and right corners of the eyes, left and right tragus, left and right trapezius points (neck), 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 (if known) 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 anthropometrics 1–8 are all collinear with the medial axis.

Finally, a word about the required anthropometric statis-



Figure 2. Landmark localization in frontal-view image. 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)

tics, namely the lower and upper bounds of both each anthropometric and each anthropometric ratio, and the mean and covariance of all anthropometrics. The statistics were all obtained from the seminal anthropometric survey in [5]. The details are given in [2]. Table 1 shows these statistics for some anthropometrics, with $\alpha = 1e-12$. It is important to note, however, that these statistics are **categorized by gender**, i.e. separate statistics for males and females. Consequently, the gender of the person in the image is assumed to be known, which is an important consideration for automated implementations of the method.

Table 1. Some anthropometric statistics (in cm), where SD: standard deviation, UB: upper bound, LB: lower bound. Based on the 1988 US Army survey [5] and with $\alpha = 2.9e-7$ (which under the assumption of normality means the upper and lower bounds correspond to 5 standard deviations from the mean).

	Males				Females			
	Mean	SD	LB	UB	Mean	SD	LB	UB
Body height	175.5	6.68	142.2	209.0	162.9	6.36	131.1	194.7
Head length	23.2	.88	18.8	27.6	21.76	.85	17.5	26.0
Bitragion	14.5	.60	11.5	17.5	13.64	.52	11.0	16.2
Interocular	10.2	.54	7.50	12.9	9.62	.50	7.12	12.1




Figure 3. Methodology for (random) generation of synthetic anthropometric data (namely y vectors) to be used in testing our anthropometric estimation method.



Figure 4. Sample images from our inhouse dataset.

3. Experiments and Results

3.1. The Data

We tested our anthropometric estimation method using both synthetic data and real images. The former consists of randomly generated y vectors, each of which is obtained by first randomly generating a vector x from the multivariate Gaussian distribution of male or female anthropometrics, then computing the corresponding pairwise ratios (Fig. 3).

The real images consist of a set of high-resolution (4368x2912) images captured in-house using a Canon 28-200mm EOS camera (Fig. 4). Each image is a **full-body shot** of one person. The dataset contains a total 108 images, with 96 frontal-view shots. Also, the dataset comprises 27 different adults, 7 females and 20 males, spanning various ethnicities: caucasian, chinese, indian, and african. 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.

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.

The body height estimation error is computed as follows: $\text{estimated value} - \text{true value}$. For the method parameters, we use $\lambda = 0.01$ and $\alpha = 2.9e-7$ (i.e. upper and lower bounds correspond to 5 standard deviations from the mean). In order to investigate the effect of choice of the anthropometric set on estimation error, we tested the method with various different sets. The results are given below in Figures 5–8 in two forms: (a) the cumulative distribution of absolute estimation error, and (b) the distribution of estimation error (as a *boxplot*). Furthermore, for comparison purposes, we have included in each Figure the performance of the baseline algorithm, i.e. the one that estimates the unknown body height as the population mean (175.5cm for males and 162.9cm for females—See Table 1).


As far as estimation from whole-body images, our method is significantly better than the baseline algorithm for synthetic data, but only slightly better for real data. Since the synthetic vectors are uncorrupted by noise, the synthetic estimation error actually corresponds to the method’s *model error*. This suggests that the method is quite sensitive to input error in the real data, which mostly comes from landmark localization and vanishing point/line estimation. Furthermore, generally speaking, the method’s performance improves when more anthropometrics are used in the estimation. This especially obvious for synthetic data. As to estimation from upper-body images, our method is no better than the baseline algorithm, both for synthetic and real data! This suggests that the main source of error lies in the linear prediction model that we currently use.

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 method was tested both synthetic data and real images. As far as whole-body images, performance is slightly better than the baseline algorithm (that randomly guesses the population mean), but inferior to that of other well-known single-view metrology techniques such as [9]. This is expected since our method uses less input information. Furthermore, since quite a bit of estimation error is caused by input error in the localization of landmarks and computation of vanishing points, this error can be reduced by estimating from multiple images or a video sequence. As for upper-body images, our method is currently no better than the baseline algorithm. However, there is room for improvement by using




(a)



(a)



(b)



(b)

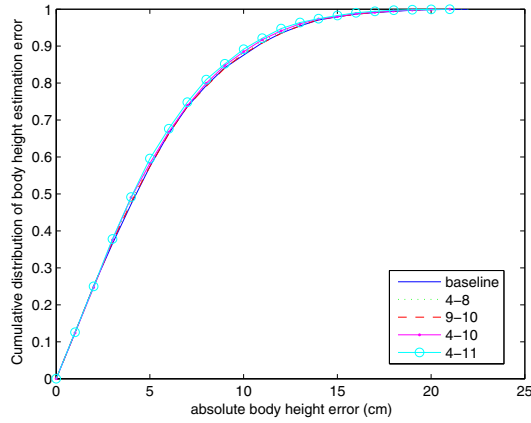
Figure 5. Synthetic data results for body height estimation from a whole-body image.

Figure 6. Real data results for body height estimation from a whole-body image.


a better prediction model than the naive linear model we currently use. In future work, we plan to: (i) explore other prediction models, particularly non-linear ones, for estimating body height from upper-body images; (ii) extend the method to estimation from multiple images; and (iii) investigate the effect of camera viewpoint (i.e. camera angle with respect to person) on estimation accuracy.

References


- [1] 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.
- [2] C. BenAbdelkader and Y. Yacoob. *Computational Forensics*, volume 1, chapter Statistical Estimation of Human Anthropometry from a Single Uncalibrated Image, pages 200–220. Springer Verlag, 2008.
- [3] A. Bovyryn and K. Rodyushkin. Human height prediction and roads estimation. In *IEEE Conference on Advanced Video and Signal-Based Surveillance*, 2005.
- [4] S. Bramble, D. Compton, and L. Klasen. Forensic image analysis. In *13th Interpol Forensic Science Symposium*, 2001.
- [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.
- [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.
- [7] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. In *CVPR*, pages 434–442, Sept. 1999.
- [8] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. *IJCV*, 40(2):123–148, Nov. 2000.




(a)



(a)



(b)



(b)

Figure 7. Synthetic data results for body height estimation from an upper-body image.

Figure 8. Real data results for body height estimation from an upper-body image.

- [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.
- [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.
- [11] L. G. Farkas, editor. *Anthropometry of the Head and Face*. Raven Press, second edition, 1994.
- [12] F. Guo. *Calibration and Metrology Using Still and Video Images*. PhD thesis, University of Maryland, College Park, 2007.
- [13] 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.
- [14] R. Hartley and A. Zisserman. *Multiple View Geometry in*

Computer Vision. Cambridge University Press, second edition, 2004.

- [15] 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.
- [16] C. Madden and M. Piccardi. Height measurement as a session-based biometric. In *Image and Vision Computing New Zealand*, 2005.
- [17] S. Pheasant. *Bodyspace: Anthropometry, Ergonomics and the Design of Work*. Taylor and Francis, second edition, 1996.
- [18] N. Saitoh, K. Kurosawa, and K. Kuroki. A study on height measurement from a single view. In *ICIP*, 1999.
- [19] G. Strang. *Introduction to Linear Algebra*. Wellesley-Cambridge Press, third edition, 2003.