Show me your face and I will tell you your height, weight and body mass index

Antitza Dantcheva
INRIA
Sophia Antipolis, France
Email: antitza.dantcheva@inria.fr

Francois Bremond
INRIA
Sophia Antipolis, France
Email: francois.bremond@inria.fr

Piotr Bilinski
University of Oxford
United Kingdom
Email: piotrb@robots.ox.ac.uk

Abstract—**Body height, weight, as well as the associated and composite body mass index (BMI) are human attributes of pertinence due to their use in a number of applications including surveillance, re-identification, image retrieval systems, as well as healthcare. Previous work on automated estimation of height, weight and BMI has predominantly focused on 2D and 3D full-*body* images and videos. Little attention has been given to the use of face for estimating such traits. Motivated by the above, we here explore the possibility of estimating height, weight and BMI from single-shot facial images by proposing a regression method based on the 50-layers ResNet-architecture. In addition, we present a novel dataset consisting of 1026 subjects and show results, which suggest that facial images contain discriminatory information pertaining to height, weight and BMI, comparable to that of body-images and videos. Finally, we perform a gender-based analysis of the prediction of height, weight and BMI.**

I. INTRODUCTION

The human face exhibits information pertaining to identity, a person's disposition, demeanor, as well as to attributes such as gender, age and ethnicity. From the perspective of biometrics, emphasis has predominantly been placed on facial recognition. More recently *attributes* or *soft biometrics* [1], [2], [3], [**?**], [**?**] such as gender, age, height and weight have gained popularity due to their semantic interpretation, *i.e.,* they can provide a description that can be readily understood by humans; for example the description "young, female, tall". Limited attention has been given to the connection between the human face and body characteristics such as body height and weight and even less so on the automatic extraction of such.

Estimating body height, weight and the associated BMI is warranted for several reasons. Firstly height and weight are attributes frequently used in surveillance, forensics, as well as re-identification applications and image retrieval systems [4]. Secondly, height and weight are primary and obvious attributes used by humans to verbally describe a person often used in police reports, unlike traditional biometrics which may be insufficient, as this was argued, for example, by Klontz and Jain [5] in the case of the 2013 Boston bombings. Thirdly, body weight and height have been proposed as soft biometric traits in automated biometric systems [**?**]. Fourthly, weight is a pertinent indicator for health and excessive weight has been associated to obesity, diabetes, and cardiovascular diseases. In this context, the presented method contributes to the current trend of *image-based automated self diagnostic*. Finally, this work can promote research conducted in psychology related to human metrology [6], [7], [8].

A. Contributions

Motivated by the above, in this work we (a) propose a new ResNet-based method for estimating body weight, height and BMI from the face. In order to conduct this study, we (b) assemble a new dataset of facial images, annotated for gender, body height, weight and BMI. In addition, (c) we provide an analysis on the impact of gender on the above estimation.

The rest of the paper is organized as follow. In Section II we review work on methods used for the estimation of height, weight and BMI. Section III introduces the proposed face-based height, weight and BMI estimation algorithm. Datasets and experimental protocol are described in Section IV. Section V presents experiments validating the effectiveness of the proposed method and a discussion thereof. Finally, Section VI concludes the paper.

II. RELATED WORK

To the best of our knowledge, the only work related to automated face-based estimation of *BMI* is a study by Wen and Guo [9], based on the MORPH-II dataset, which obtained mean absolute errors (MAEs) for BMI in the range from $2.65-4.29$ for different ethnic categories. The study explored handcrafted features for BMI-estimation and specifically in the method the face was detected, normalized, and an active shape model was fitted, based on which, geometry and ratio features were extracted (cheekbone to jaw width, width to upper facial height ratio, perimeter to area ratio, eye size, lower face to face height ratio, face width to lower face height ratio and mean of eyebrow height), normalized and finally subjected to support vector regression. We note that the BMI-annotation of MORPH-II has not been made public.

Body weight and BMI have received an increased attention in the recent years. This relates to different trends that have appeared, such as for example the fact that more than two-thirds of U.S. adults are now overweight or obese [10], and the fact that obesity - which is now recognized as a disease by the American Medical Association[1] - accounts for nearly one

[1]www.ama-assn.org/go/obesity

out of every 10 American deaths. Specifically an increased body weight and thus BMI, is associated with many health problems such as diabetes, heart disease, certain cancers , strokes, osteoarthritis and gallstones, all of which cost the United States 223 billion per year[2]. These motivate the use of automated extraction of body weight and of BMI - computed from body weight and body height as:

$$BMI = \frac{weight[kg]}{height^2[m^2]} = \frac{weight[lb] \cdot 703}{height^2[in^2]}, \tag{1}$$

which further has the advantage of being contactless, low-cost, unobtrusive, and unconstrained[3].

Body weight has been estimated from manually measured *body-anthropometric data* [11] by constructing a weight prediction metric using multiple regression analysis. Related experiments on the NHANES dataset[4] - which contains data of 28,000 subjects - showed that weight could be estimated in 93% of the test data with an error of less than 10%. In addition, an automated and Kinect based weight estimation [12] from full-body shots obtained weight estimation error of 4% in a small dataset of 15 subjects. Furthermore a related work for weight estimation by Labati et al. [13] studied frame sequences representing walking subjects. Under specific video acquisition from calibrated cameras, body height, body volume (estimated by ellipses passing through the silhouettes), body shape and walking direction were extracted, to train a neural network towards body weight estimation. The mean errors ranged from -2.38 kg to 0.02 kg, and they corresponded to a private dataset of 20 subjects.

Human estimation of body weight from face and body [14] has shown to be rather imprecise, with 47% of estimates at least 10% different and 19% of estimates at least 20% different from the measured values. The majority of height estimates were within 10% of the measured values. Similar results were obtained by a study by Coe et al. [15], where four observers estimated the weight and height of 38 patients who were lying covered on operating theatre trolleys. In addition, drug dosages and infusion rates in adults are usually based on body weight, as are nutrition prescriptions [16], [17]. Nomograms, such as the Cockcroft-Gault and the Harris-Benedict formulas, as well as intensive care unit (ICU) scoring systems, such as SOFA (Sequential Organ Failure Assessment), rely on knowledge of the patients weight [18], [19], [20].

On a different, but related note concerning correlation of face and body metrology the work by Adjeroh et al. [8] used the manual 1D-measurements of the 3D CAESAR database[5] to show a high degree of correlation clusters between face and body measurements. Utilizing the existing correlation, gender could be estimated with a success rate of 100% (this corresponded to 44 measurements), and body weight could be estimated with a success rate of 92.5%. Furthermore Criminisi et al. [21] computed 3D affine measurements from single perspective views, where affine scene structure could be determined from the image, without knowledge of the camera's internal calibration (*e.g.*, focal length), nor of the explicit human pose.

Height is also required for calculations of ideal body weight or adjusted body weight [22], [23], and height or weight is used to calculate and communicate ventilator settings [24].

Both weight and height are required to estimate body surface area, cardiac index and body mass index [22], [25]. In addition, application of many commonly used clinical guidelines and evidence-based practices require the knowledge of height and weight. We here note that while generally, patients can be weighed and measured or questioned, in some cases (*i.e.,* elderly trauma patients or intensive care patients) measuring or questioning is not an option due to decreasing mental ability, as well as immobility or unconsciousness. Finally, we note that the failure to obtain and record weight and height has been highlighted in the National Confidential Enquiry into Peri-operative Deaths (NCEPOD)[26] reports, concluding that accurate estimation or measurement of height and weight of critically ill patients is a pertinent prereposite to optimum clinical care.

A. Convolutional Neural Networks

Deep Convolutional Neural Networks have significantly advanced attribute recognition, where existing architectures such as AlexNet [27], VGG [28] and Inception [29] have been employed and adapted. One finding has been that when pre-training with massive face images, high-level hidden neurons automatically learn semantic concepts and that such concepts are significantly enriched after fine-tuning with attribute tags. Further, a deep model pre-trained for face recognition, implicitly learns attributes.

III. Proposed Approach

Motivated by the above, we propose here a method for face-based estimation of height, weight and BMI based on the ResNet-50 architecture.

We pose the problems of height, weight and BMI estimation as 3 separate regression pattern classification problems. We employ face detection prior to the ResNet as an attention mechanicsm. While deep neural networks (DNNs) tackle challenging settings of *attributes in the wild* that encompass complex face variations such as poses, lightings, and occlusions reasonably well and generally no face detection is required, we ensure by detecting the face, that *no body information is considered* in our study and hence we analyze the *face-based* height, weight and BMI estimation.

In what follows, we proceed to describe briefly the pre-processing step of face detection and the employed ResNet-architecture.

[2]http://www.health.harvard.edu/newsletters/harvard_mens_health_watch/2012/February

[3]Extracting BMI or body weight offers also advantages, even in forensics, as these measures can be sensed at a distance and are relatively distinctive and permanent

[4]http://www.cdc.gov/nchs/nhanes.htm

[5]http://www.hec.afrl.af.mil/cardlab/CAESAR/index.html

A. Face Detection

There exist a large number of face detection algorithms, based on a large number of features and implementations. For our employed datasets the classic Viola-Jones algorithm [30] performed very well (100% correctly detected faces) and we hence incorporate it in the proposed algorithm.

We resize the detected faces to 256×256, and provide these as input of ResNet.

B. Residual Neural Networks

We employ in this work a residual network, the 50-layer ResNet [31] architecture. While training deep neural networks bared challenges such as exploding / vanishing gradients and degradation, residual links met the challenge and have shown to surpass other techniques such as initialization strategies, better optimizers, skip connections, knowledge transfer, as well as layer-wise training enabling training of deeper neural networks. The incorporated residual connections have been inherently necessary for training very deep convolutional networks. ResNets have significantly progressed accuracy of object classification, object detection and segmentation, while improving the training speed.

Residual blocks with identity mapping can be represented by the following

$$\mathbf{x_{l+1}} = \mathbf{x_l} + \mathcal{F}(\mathbf{x_l}, \mathcal{W}_l), \tag{2}$$

where $\mathbf{x_{l+1}}$ and $\mathbf{x_l}$ are input and output of the *l*-th unit in the network, $\mathcal{F}$ is a residual function and $\mathcal{W}_l$ are parameters of the block. The characteristic of residual networks lies in their architecture based on sequentially stacked residual blocks. We use basic blocks with two consecutive 3×3 convolutions with batch normalization and ReLU preceding convolution.

Layer Name	
Conv1	7×7, 64, stride 2
Conv2_x	3× 3 max *pool*, *stride* 2 $\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$
Conv3_x	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 4$
Conv4_x	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 1024 \end{bmatrix} \times 6$
Conv5_x	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512 \\ 1 \times 1, 2048 \end{bmatrix} \times 3$
	Average pool, 1000-d fc, smooth l1
FLOPs	3.8×10^9

Fig. 1. 50-layer ResNet architecture used in this work.

We use the original $conv - BN - ReLU$ for order of batch normalization, activation and convolution in residual blocks.

C. Implementation Details

We replaced the initial $1000{-}d$ fc of ResNet with a $1{-}d$ fc and then Smooth L1 Loss in order to cater regression.

Experimental Setup: We train our network end-to-end with the classical back-propagation algorithm using a single NVIDIA GeForce GTX TITAN X with 12 GB memory. We implement our approach in the Caffe deep learning framework [**?**].

A 224×224 crop is randomly sampled from each image or its horizontal flip, with the per-pixel mean subtracted following the work of He et al. [31]. In addition, we initialize the weights using those from the ResNet-50 model [31], which were pre-trained on the ImageNet Large-Scale Visual Recognition Challenge (ILSVRC) 2014 [32] (1000-class ImageNet classification task), and then we fine-tune in turn the finer strides on the $VIP\_attribute$ dataset. We use the stochastic gradient descent (SGD) algorithm with a mini-batch size of 16. The learning rate starts from $1e^{-5}$ and then we use the "inv" learning rate decay policy, which is implemented as follows:

$$base_{lr} * (1 + gamma * iter)^{(-power)}, \tag{3}$$

where gamma is 0.001 and power is 0.75. We use a weight decay of 0.0005 and a momentum of 0.9. The models are trained for up to $400,000$ iterations.

IV. Experiments

A. Dataset

For this study, we assembled a dataset, which we refer to as the $VIP\_attribute$ dataset, consisting of 1026 subjects, specifically 513 female and 513 male celebrities (mainly actors, singers and athletes) collected from the WWW.The images are mainly frontal images. Co-variates include illumination, expression, image quality and resolution. Further challenging are beautification (photoshop) of the images, as well as the presence of makeup, plastic surgery, beard and mustache. We obtained annotations related to the subjects' body weight and height available on websites such as www.celebheights.com, www.howtallis.org and celebsize.com, and proceeded to calculate the associated BMI. Figure 2 illustrates example images of the $VIP\_attribute$ dataset.

Table I reports the statistics of gender, height, weight and BMI associated to the dataset. The related histograms are depicted in Figure 3. We note that the annotation might not be fully accurate, *e.g.,* weight fluctuations throughout time or inaccurate self-report. Images from the $VIP\_attribute$ dataset are available in the authors webpage: http://www.antitza.com/VIP_attribute-dataset.html.

B. Correlation between height, weight, BMI and gender

Given two vectors $Y = y_1, y_2, ..., y_n$ and $Z = z_1, z_2, ..., z_n$ the Pearson correlation coefficient is computed as

$$\rho_{Y,Z} = \frac{cov(Y,Z)}{\sigma_Y \sigma_Z} = \frac{E(Y\overline{Y})(Z\overline{Z})}{\sigma_Y \sigma_Z}, \tag{4}$$

where σ_Y and σ_Z are the respective standard deviations for Y and Z. The coefficient ranges between -1 and 1, with the two extreme points being obtained when the variables are maximally linearly related. We compute the correlation coefficients between each pair of {gender height, weight,

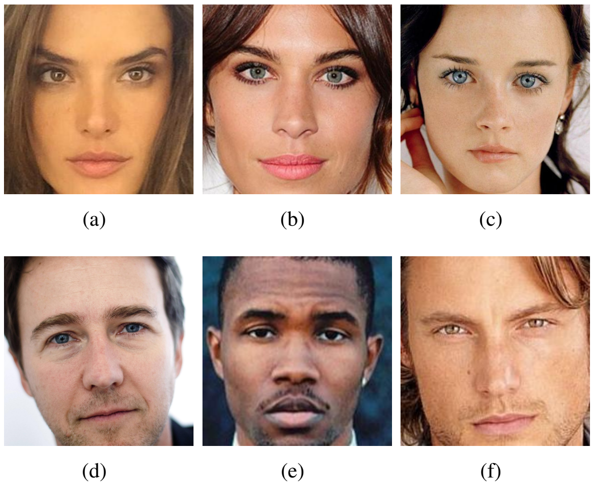

Fig. 2. Examples of six subjects in the $VIP\_attribute$ dataset assembled by the authors.

TABLE I
CHARACTERISTICS OF $VIP\_attribute\_dataset$. MEAN AND STANDARD DEVIATIONS (σ) PERTAINED TO WEIGHT, HEIGHT AND BODY MASS INDEX FOR FEMALES AND MALES.

	Female	**Male**
Subjects	513	513
Height [m]	1.67 ($\sigma = 0.09$)	1.85 ($\sigma = 0.14$)
Weight [cm]	58.34 ($\sigma = 11.02$)	86.93 ($\sigma = 21.00$)
Body mass index	20.87 ($\sigma = 3.71$)	25.21 ($\sigma = 3.57$)

gender}, see Table II. We observe naturally a large correlation between BMI and weight, lesser with height which is due to the fact that ρ_{YZ} represents a *linear* correlation between Y and Z. Interestingly body weight and body height have intrinsically a large correlation. In agreement with previous studies [33], height, weight and hence BMI are strongly correlated with gender.

TABLE II
PEARSON'S CORRELATION COEFFICIENT FOR GENDER, HEIGHT, WEIGHT AND BMI IN THE $VIP\_attribute$ DATASET.

	gender	**height**	**weight**	**BMI**
gender	1	0.59	0.65	0.51
height	0.59	1	0.79	0.27
weight	0.65	0.79	1	0.78
BMI	0.51	0.27	0.78	1

V. RESULTS

In this section, we evaluate the performance of the proposed height, weight and BMI estimating methods on the $VIP\_attribute$ dataset. As accuracy measures we report the Mean Absolute Error (MAE) and the Pearson's correlation coefficient ρ for each estimation related to height, weight and BMI. We compute the MAEs as the average of the absolute errors between the estimated height, weight and BMI and the respective ground truth. $MAE = \frac{1}{N}\sum_{k=1}^{N}|\hat{b_k} - b_k|$, where b_k is the ground truth for image k, $hatb_k$ is the associated estimated value (height, weight or BMI) and N is the total number of test images. MAEs has been often used as performance measure in age estimation [2].

Higher errors generally related to atypical cases such as the cases of 'under weight' and 'obese' subjects, as well as to 'short' and 'very tall' subjects, due to the limited number of training samples.

A. Height

As presented in Table III, when estimating height from faces, MAEs [m] in the range of $0.077 + 0.005$ were obtained. Given the statistics shown in Table I, this corresponds to $< 5\%$ of the female, male and hence overall average height, comparatively similar to human estimation [14], which reportedly lay within the the $< 10\%$ range. The related ρs range between 0.57 and 0.60. Notable is that while MAEs pertained to females and males are smaller than the MAE of all subjects combined, the related correlation rates are higher for the latter. Overall we cannot observe a gender-bias in height-estimation.

TABLE III
RESULTS OF HEIGHT ESTIMATION [M]. ABBREVIATIONS USED: MAE: MEAN ABSOLUTE ERROR, ρ: PEARSON'S CORRELATION COEFFICIENT.

	Male	**Female**	**All**
MAE [m]	0.077	0.078	0.082
ρ	0.57	0.57	0.60

B. Weight

We report the regression accuracy related to weight in Table IV. We observe MAEs between $8.06kg$ (females) and $8.51kg$ (all subjects). The lower MAE of females is intuitive, since females have generally a lower body weight (see Table I and Figure 3). It constitutes though about 14% of the mean weight. This is marginally larger than the weight estimation, given several precise body-measurements [11], which constituted in reportedly less than 10% in 93% of the cases. Kinect-aided weight estimation based on body images has shown though a significantly better performance [12], [13]. The associated human accuracy is though rather imprecise too, with 47% of estimates at least 10% different and 19% of estimates at least 20% different from the measured values [14]. On the other hand for females we have a very high correlation rate of $\rho > 0.78$.

We envision, that specifically what contributes to the error in weight estimation are factors including: imprecise weight annotation, due to weight-fluctuations in time, as well as which might have been purposely augmented in favor of the subjects, face alterations such as (facial) plastic surgery, facial makeup and beautification of images (*e.g.,* by Photoshop).

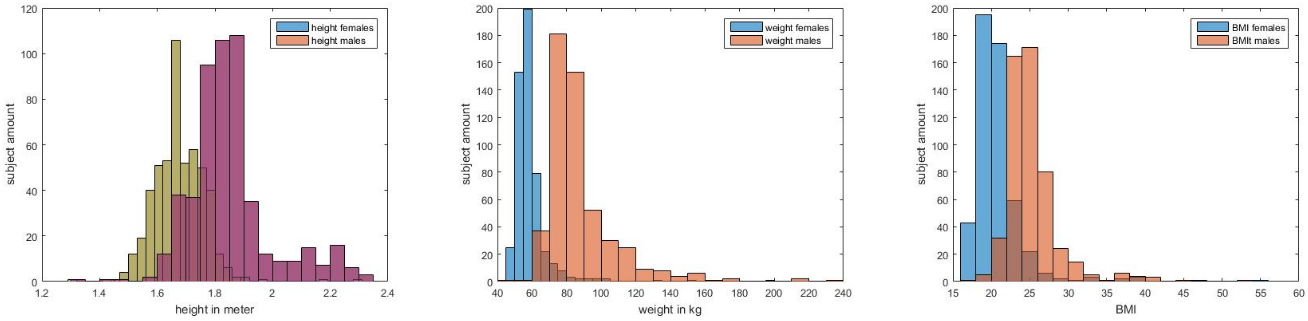

(a) (b) (c)

Fig. 3. Histograms of (a) height, (b) weight and (c) BMI in the $VIP\_attribute$ dataset assembled by the authors.

TABLE IV
RESULTS OF WEIGHT ESTIMATION [KG]. ABBREVIATIONS USED: MAE: MEAN ABSOLUTE ERROR, ρ: PEARSON'S CORRELATION COEFFICIENT.

	Male	**Female**	**All**
MAE [kg]	8.25	8.06	8.51
ρ	0.76	0.78	0.75

C. BMI

We report the BMI-regression accuracy in Table V. The MAEs related to BMI are in the range of $2.3 + 0.06$, which is relatively small, given the absolute range of BMI, predominantly between $15 - 35$, illustrated in Figure 3 (c). We outperform the results of Wen and Guo *et al.* [9], which were $2.65 - 4.29$. The fluctuations between female, male and overall estimation with respect to MAEs and ρ are minimal, which indicates that BMI-estimation is not gender-biased.

TABLE V
RESULTS OF BMI ESTIMATION. ABBREVIATIONS USED: MAE: MEAN ABSOLUTE ERROR, ρ: PEARSON'S CORRELATION COEFFICIENT.

	Male	**Female**	**All**
MAE	2.32	2.30	2.36
ρ	0.55	0.55	0.55

VI. SUMMARY

This paper presented a novel approach for estimating height, weight and BMI *from single-shot facial images*, based on ResNet-50. Experiments conducted on a new dataset, which we have made publicly available, resulted in promising correlation accuracies of up to $\rho = 0.78$ for female weight estimation and mean absolute errors of 2.3 for female BMI-estimation. We did not observe a significant gender-bias in estimating height, weight and BMI. However, more work is necessary in this regard. Future work will involve the additional study of age and ethnicity in order to improve utilization of facial appearance for height, weight and BMI estimation. The height, weight and BMI estimator was motivated by the current need for self-diagnostic tools for remote healthcare, as well as for soft biometrics categorization in security applications.

REFERENCES

[1] A. K. Jain, S. C. Dass, and K. Nandakumar, "Soft biometric traits for personal recognition systems," in *Proc. of ICBA*, 2004.
[2] A. Dantcheva, P. Elia, and A. Ross, "What else does your biometric data reveal? a survey on soft biometrics," *IEEE Transactions on Information Forensics and Security*, pp. 1–26, 2015.
[3] M. S. Nixon, P. L. Correia, K. Nasrollahi, T. B. Moeslund, A. Hadid, and M. Tistarelli, "On soft biometrics," *Pattern Recognition Letters*, vol. 68, pp. 218–230, 2015.
[4] J. Bekios-Calfa, J. M. Buenaposada, and L. Baumela, "Robust gender recognition by exploiting facial attributes dependencies," *PRL*, 2013.
[5] J. C. Klontz and A. K. Jain, "A case study on unconstrained facial recognition using the boston marathon bombings suspects," *Technical Report*, no. MSU-CSE-13-4, 2013.
[6] V. Coetzee, J. Chen, D. I. Perrett, and I. D. Stephen, "Deciphering faces: Quantifiable visual cues to weight," *Perception*, vol. 39, no. 1, pp. 51–61, 2010.
[7] D. D. Pham, J.-H. Do, B. Ku, H. J. Lee, H. Kim, and J. Y. Kim, "Body mass index and facial cues in sasang typology for young and elderly persons," *Evidence-Based Complementary and Alternative Medicine*, vol. 2011, 2011.
[8] D. Adjeroh, D. Cao, M. Piccirilli, and A. Ross, "Predictability and correlation in human metrology," in *Proc. of the IEEE WIFS*, 2010.
[9] L. Wen and G.-D. Guo, "A computational approach to body mass index prediction from face images," *Image and Vision Computing*, vol. 31, no. 5, pp. 392–400, 2013.
[10] E. A. Yates, A. K. Macpherson, and J. L. Kuk, "Secular trends in the diagnosis and treatment of obesity among us adults in the primary care setting," *Obesity*, vol. 20, no. 9, p. 19091914, 2012.
[11] C. Velardo and J.-L. Dugelay, "Weight estimation from visual body appearance," in *Proc. of IEEE BTAS*, 2010.
[12] C. Velardo, J. L. Dugelay, M. Paleari, and P. Ariano, "Building the space scale or how to weigh a person with no gravity," in *Proc. of ESPA*, 2012.
[13] R. D. Labati, A. Genovese, V. Piuri, and F. Scotti, "Weight estimation from frame sequences using computational intelligence techniques," in *Proc. of CIMSA*, 2012.
[14] R. Bloomfield, E. Steel, G. MacLennan, and D. W. Noble, "Accuracy of weight and height estimation in an intensive care unit: Implications for clinical practice and research," *Critical care medicine*, vol. 34, no. 8, pp. 2153–2157, 2006.
[15] T. R. Coe, M. Halkes, K. Houghton, and D. Jefferson, "The accuracy of visual estimation of weight and height in pre-operative supine patients," *Anaesthesia*, vol. 54, p. 582598, 1999.
[16] D. Morgan and K. Bray, "Lean body mass as a predictor of drug dosage: Implications for drug therapy," *Clin Pharmacokinet*, vol. 26, p. 292307, 1994.
[17] P. Jolliet, C. Pichard, and R. Biolo, "Enteral nutrition in intensive care patients: A practical approach," *Intensive Care Med*, vol. 24, p. 848859, 1998.
[18] D. Cockcroft and M. Gault, "Prediction of creatinine clearance from serum creatinine," *Nephron*, vol. 16, p. 3141, 1976.

[19] F. Simpson and G. Doig, "Parenteral vs. enteral nutrition in the critically ill patient: A metaanalysis of trials using the intention to treat principle," *Intensive Care Med*, vol. 31, p. 1223, 2005.

[20] J. Vincent, R. Moreno, and J. e. a. Takala, "The sofa sepsis-related organ failure assessment score to describe organ dysfunction failure," *Intensive Care Med*, vol. 22, p. 707710, 1996.

[21] A. Criminisi, I. Reid, and A. Zisserman, "Single view metrology," *International Journal of Computer Vision (IJCV)*, vol. 40, no. 2, pp. 123–148, 2000.

[22] D. Carney and M. Meguid, "Current concepts in nutritional assessment," *Arch Surg*, vol. 137, p. 4244, 2002.

[23] E. BL, "Dosing of medications in morbidly obese patients in the intensive care unit setting," *Intensive Care Med*, vol. 30, p. 1832, 2004.

[24] P. Eichacker, E. Gerstenberger, and S. e. a. Banks, "Meta-analysis of acute lung injury and acute respiratory distress syndrome trials testing low tidal volumes," *Am J Respir Crit Care Med*, vol. 166, p. 15101514, 2002.

[25] S. Gibson and A. Numa, "The importance of metabolic rate and the folly of body surface area calculations," *Anaesthesia*, vol. 58, p. 5055, 2003.

[26] K. Callum, N. Carr, A. Gray *et al.*, "National confidential enquiry into perioperative deaths," *Extremes of age: the 1999 report of the National Confidential Enquiry into Perioperative Deaths: data collection period*, vol. 1, 1997.

[27] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 3730–3738.

[28] K. He, Y. Fu, and X. Xue, "A jointly learned deep architecture for facial attribute analysis and face detection in the wild," *arXiv preprint arXiv:1707.08705*, 2017.

[29] M. S. Sarfraz, A. Schumann, Y. Wang, and R. Stiefelhagen, "Deep view-sensitive pedestrian attribute inference in an end-to-end model," *arXiv preprint arXiv:1707.06089*, 2017.

[30] P. Viola and M. Jones, "Robust real-time face detection," in *Proceedings of IEEE ICCV*, 2001.

[31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[32] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.

[33] D. Cao, C. Chen, M. Piccirilli, D. Adjeroh, T. Bourlai, and A. Ross, "Can facial metrology predict gender?" in *Proc. of IJCB*, 2011.