

Evaluation of Features from RGB-D Data for Human Body Weight Estimation

Christian Pfitzner* Stefan May* Andreas Nüchter**

* *Nuremberg Institute of Technology Georg Simon Ohm,
Kesslerplatz 12, 90489 Nuremberg, Germany
(e-mail: christian.pfitzner@th-nuernberg.de).*

** *Julius-Maximilians-Universität Würzburg,
Am Hubland, 97074 Würzburg, Germany*

Abstract: Body weight is a crucial parameter when it comes to drug or radiation dosing. In case of emergency treatment time is short so that physicians estimate the body weight by the visual appearance of a patient. Further, visual body weight estimation might be a feature for person identification. This paper presents the anthropometric feature extraction from RGB-D sensor data, recorded from frontal view. The features are forwarded to an artificial neural network for weight estimation. Experiments with 233 people demonstrate the capability of different features for body weight estimation. To prove robustness against sensor modalities, a structured light sensor is used, as well as a time-of-flight sensor. An additional experiment including temperature features from a thermal camera improves the body weight estimation beyond.

Keywords: human body weight, anthropometric features, cognition, rgb-d, thermal camera, neural network

1. INTRODUCTION

Body weight might be one of the best indicators to check health status: Under or overweight can be an evidence for sickness. Further, body weight estimation is crucial for some clinical scenarios: A physician treating a patient with an acute ischemic stroke has to estimate the body weight to adapt the drug dosage. The medication to solve a blood clot has to be given within a narrow time window of three hours after appearing of the first symptoms of a stroke. Under this time pressure it is still state of the art for physicians to estimate a patient's body weight by a visual guess. Several studies demonstrated that the physicians and nursing staff can hardly estimate someone's weight sufficiently for drug dosing (Fernandes et al., 1999; Breuer et al., 2010). In case of elderly patients, dementia might interfere with the patient's self estimation. In addition, not everybody measures themselves regularly and steps on a scale. Bed scales are not available in every trauma room and in addition high errors in weight measurement appear if the wrong tare weight of the bed is taken. There exist several anthropometric body weight estimation methods for clinical usage which give a fast but rough estimation of body weight based on the measurement of lengths and circumferences of the human body with a measuring tape (Buckley et al., 2012).

An other application where the visual body weight can play an important role is identification: Some soft biometrics, e.g. color or length of the hair, can be changed quickly. In contrast, body weight can not be changed immediately. Although it can be disguised, e.g. with thick clothes or a

corsage, body weight can help in the identification of a single person (Ailisto et al., 2006). The body weight and the constitutional type of a person are visible from distance in contrast to close-up visible soft biometrics, e.g. color of the eyes (Dantcheva et al., 2010).

The contributions of this paper are as follows: First, the previous approach by Pfitzner et al. (2016) is extended with additional geometric features. Second, all features are analyzed towards their capability for body weight estimation, having an RGB-D dataset provided. Finally, extended experiments demonstrate the improvement of the algorithm, with an extension for visual body weight estimation without volume prediction. Data from a thermal camera can improve the outcome in body weight estimation further, as well as knowledge of the gender. Additionally, the data from experiments is provided online, including ground truth values for body weight.

The paper is structured as follows: Related work provides an overlook about body weight estimation based on visual sensors. In the following section the algorithm for feature extraction is described. Further, experiments demonstrate the capability of different feature groups to estimate body weight in comparison to related work. These experiments rely on 233 datasets and demonstrate the impact of different features to provide an estimation of human body weight. The approach is verified with data from a structured light sensor (Microsoft Kinect) – and a time-of-flight sensor (Microsoft Kinect One). Finally, the last section gives a conclusion with an outlook to future work.

* This work is funded by the Federal Ministry of Education and Research in Germany. Project funding reference number: 03FH040PX3.

2. RELATED WORK

Since the release of the Microsoft Kinect camera, 3D perception got a boost due to the fact of being a low-cost consumer sensor. Several publications exist with this sensor concerning the extraction of anthropometric features. The position of the human skeleton was essentially important to use the camera in gaming and multimedia applications. The extraction and labeling of skeleton features was provided by Shotton et al. (2013): In 900,000 synthetic images from the Kinect camera, a labeling for body parts was generated with the help of a random decision tree. Further, people can be identified from skeleton and by their gait (Gabel et al., 2012).

Also anthropometric features from depth data can be found in related work: Pirker et al. (2010) estimated human body volume in clinical environment by eight stereo cameras around a stretcher and bioelectrical impedance analysis. With a 3D reconstruction the volume of a frontal surface can be calculated towards the medical stretcher which has been modeled as a plane.

Velardo and Dugelay (2012) extrapolated anthropometric features from body silhouette. The silhouette is proved by a RGB-D sensor. The extracted features, e.g. height, waist circumference, and the length and circumference of arm and leg, are used to train a statistical model.

Further, Robinson and Parkinson (2013) demonstrated the extraction of anthropometric features from different poses: Their approach demonstrated that raw data from Kinect camera can produce a rough estimation of anthropometric features, due to sensor noise and bias in depth measurement. Additionally, anthropometric features are hard to provide with sufficient precision via an optical sensor: Even thin clothes might confuse the extraction of a circumference of a body part, e.g. the waist circumference. Depending on the underlying calculation of body weight these errors can have a high impact in error.

Nguyen et al. (2014) developed a method to predict body weight by a side view feature and a support vector regression model. Separating datasets by gender their approach reached an average error of 4.62 kg for females and 5.59 kg for males. Finally, they compared the body weight estimation by the algorithm in contrast to visual estimation.

Pfzner et al. (2015) demonstrated a body weight estimation by volume extraction from RGB-D data. The presented algorithm provided an accuracy of 79% for a cumulative error of $\pm 10\%$. The approach was tested with 110 patients from trauma room, focusing on body weight estimation for stroke patients. Due to uncertainties in volume estimation this approach had outliers up to 32%. Compared to an physician's estimation this approach is already more suitable for drug dosing.

Pfzner et al. (2016) provide a feature-based body weight estimation tested with 69 patients from clinical environment. Feature extraction was improved by adding temperature information from a thermal imaging device to ease segmentation of the patient. Further, an artificial neural network (ANN) was implemented and fed with additional geometric features from eigenvalues. These features rely

on robust features which hardly impair with clothes or a precise measuring pose of the person. Finally, this approach reached nearly 90% for body weight estimation within a range of $\pm 10\%$ towards ground truth body weight. Further, the range of outliers was reduced towards less than $\pm 18\%$ from ground truth.

The here presented approach extends the work of Pfzner et al. (2016) with additional features from RGB-D and thermal camera, as well as an analysis in correlation of robust features for body weight estimation.

3. APPROACH

The data for feature extraction are acquired with a Microsoft Kinect camera, a Microsoft Kinect One camera, as well as a thermal camera. Both Kinect cameras deliver color images as well as depth data. First of all the thermal camera is used to ease segmentation of a person from the environment.

For weight estimation people are placed on a stretcher, about two meters away from the sensor to be visible in the field of view. The stretcher can be adapted in height. This sensor configuration and environment was chosen on the basis of previous work, which had the focus for body weight estimation in clinical environment. There was no given pose for the test probands, however all were laying on their back with their arms beside or crossed on their stomach.

3.1 Pre-Processing

The data acquisition was done with the sensors integrated into the ceiling of a room, see Fig. 1(a). All three sensors are mounted rigid to each other to prevent a loose in extrinsic calibration. The sensors field of views is facing towards the ground. By sensor fusion a point cloud \mathcal{P} is generated containing N_c Cartesian points with color c_{rgb} and additionally the temperature t with $\mathbf{p}_j = (x \ y \ z \ c_{rgb} \ t)^T \in \mathcal{P}$. Figure 1(a) shows the scene from the sensor's view with colors overlaying by false color temperature.

Segmentation is achieved by the help of the thermal camera, filtering for a certain threshold to differ between human and environment. Additional image processing techniques like color filtering, bilateral filtering and morphological operations improve the result in segmentation. After this step the subset containing points of the person is known by $\mathcal{P}^p \subseteq \mathcal{P}$.

3.2 Feature Extraction

This section describes the extraction G of the features $\mathbf{f}_i = (f_1 \dots f_M)$ from a single point cloud \mathcal{P}_i^p . The size of a feature vector is set by M . All obtained feature vectors are stored in a feature set \mathbf{F} with the size I by

$$\mathbf{f}_i = G(\mathcal{P}_i^p) \quad \text{with} \quad \mathbf{f}_i \in \mathbf{F} \quad \text{and} \quad i = \{1, 2, \dots, I\}. \quad (1)$$

To determine which features are best for body weight estimation the correlation is investigated. Figure 4 depicts the correlation between the features itself and ground truth.

The computation of the features $f_1 - f_{13}$ is shown in previous work Pfzner et al. (2016). For the reader's

Table 1. List of features for body weight estimation $\forall \mathbf{p}_j \in \mathcal{P}^p$.

f_1	volume	v
f_2	surface	s
f_3	number of points	n
f_4	density	n/N_C
f_5	1st eigenvalue	λ_1
f_6	2nd eigenvalue	λ_2
f_7	3rd eigenvalue	λ_3
f_8	sphericity	$\lambda_3/\sum_j \lambda_i$
f_9	flatness	$2 \cdot (\lambda_2 - \lambda_3) / \sum_i \lambda_i$
f_{10}	linearity	$(\lambda_1 - \lambda_2) / \sum_i \lambda_i$
f_{11}	compactness	$\sqrt{1/n \sum_i (\mathbf{p}_j - \bar{\mathbf{p}})^2}$
f_{12}	kurtosis	$1/n \sum_j \ \mathbf{p}_j - \bar{\mathbf{p}}\ ^4$
f_{13}	alt. compactness	$\sum_j (\mathbf{p}_j - \bar{\mathbf{p}})^4 / f_9$
f_{14}	distance to person	d
f_{15}	contour length	l_c
f_{16}	contour area	a_c
f_{17}	convex hull length	l_h
f_{18}	convex hull area	a_h
f_{19}	gender	g
f_{20}	min temperature	$\min(t)$
f_{21}	max temperature	$\max(t)$
f_{22}	avg. temperature	$\text{avg}(t)$
f_{23}	ambient temperature	t_{amb}

convenience, it is presented here again in Tab. 1. A person laying on a flat surface eases volume reconstruction due to plane modeling of the back, as shown in Pfitzner et al. (2015): The volume v is extracted by RANSAC plane modeling as the back surface of a person, and a meshing approach is applied for the frontal surface of the patient, which is visible to the sensor (Fischler and Bolles, 1981). The features f_5 to f_{10} are based on the eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ from principal component analysis. Fig. 3 illustrates the geometric principle for the patient’s point cloud \mathcal{P}^p and the eigenvalues. The features f_{11} to f_{13} indicate the compactness of a point cloud after segmentation from the environment based on the centroid $\bar{\mathbf{p}} = \sum \mathbf{p}_j / n$ with $\mathbf{p}_j \in \mathcal{P}_i^p$.

The previous approach is extended with the features f_{14} – f_{23} : The distance to the patient d is added to improve body weight estimation independently from the distance between sensor and test person. From the extracted contour of a patient’s silhouette, the length l_c and the area a_c can be calculated with a convex hull around the person’s silhouette, it’s length l_h and area a_h . Figure 2 demonstrates the extracted contour and convex hull from a person’s RGB-D dataset. The gender is added as a feature to be forwarded to the ANN. Although the visual gender estimation is not integrated in the current approach, it could be implemented as shown by Linder et al. (2015). For experiments the gender is added manually.

Table 2. Survey of datasets. The dataset for training is taken from the first 25 percent of the total dataset. Both datasets contain data from Kinect camera, while *Event* also contains data from the Kinect One.

Dataset	Real Weight in kg				Gender	
	min	max	mean	σ	female	male
Hospital	48.6	129	77.8	17.1	72	55
Event	48.8	114	78.6	12.0	24	82
Total	48.6	129	79.5	15.3	137	96
Training	48.6	109.8	75.0	14.3	34	25

To compensate the error in estimation of people wearing thick clothes, features from the thermal camera are added: On one side, wearing thick clothes would result in a higher volume and a bigger surface. Due to the strong positive correlation (see Fig. 4) this has a direct effect on the outcome of weight estimation. On the other side, thick clothes insulate the body heat of a person, which can be noticed in thermal camera and a wider range for minimum and maximum temperature (t_{min} , t_{max}). The ambient temperature t_{amb} is added to the feature vector. In the applied total dataset \mathbf{F} the ambient temperature ranged from 19.1 °C to 26.8 °C. In experiments, the hot spot of a human being was generally located on the head. The here presented approach is implemented with the Point Cloud Library Rusu and Cousins (2011).

4. DATASETS

To encourage future work by other research groups the features from the recorded datasets are made public. For scientific purpose the dataset contain the values of features and ground truth body weight. The features are provided at www.provided-after-review.com.

- **Hospital \mathbf{F}_H** : This dataset contains feature values from trauma room patients from the Universitätsklinikum Erlangen, Germany. The dataset contains 127 measurements from people laying on a medical stretcher, recorded with a Microsoft Kinect. For this dataset a good distribution is achieved having people of different ages, body weights and shapes, see Tab. 1. Additionally, this dataset contains the patient’s self estimation, age, sex, as well as anthropometric features like body height, abdominal girth and waist circumference. The distance between the sensors and the probands was around 2 m.
- **Event \mathbf{F}_E** : The features from this dataset were recorded at a public event. People in this dataset were visitors of the public event. This dataset contains 106 people. Additionally, this dataset includes point clouds from Microsoft Kinect One.

All features were extracted by the previously presented algorithm.

5. SETTING FOR EXPERIMENTS

In experiment section different configurations in features will be explored. To compare the varying set-ups the errors in weight estimation are evaluated. The estimated body weight \tilde{w} is evaluated with the help of the ground

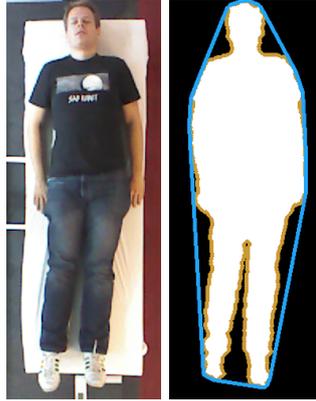


(a)



(b)

Fig. 1. Scene for body weight recorded dataset *Hospital* (a) and sensor module in the ceiling with thermal imaging camera Optris PI 400, Kinect and Kinect One. (b)

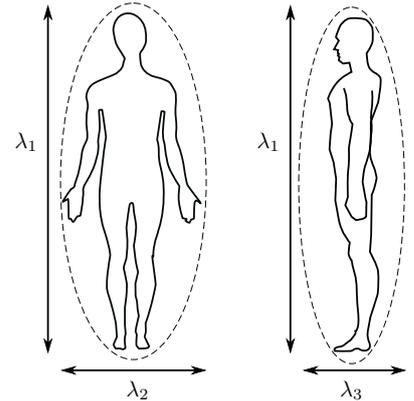


(a)



(b)

Fig. 2. Sensor's view from the ceiling with proband (a). The extracted person, its contour (orange) and convex hull (blue) are shown in (b).



(a)

(b)

Fig. 3. Principal components along the person's silhouette: The 1st eigenvalue will be set along the height of the person, the 2nd will be set as width and the 3rd eigenvalue will be set as the depth.

truth body weight \hat{w} . While the absolute error $e = \hat{w} - \tilde{w}$ is a good indicator if tested people have nearly the same body weight, the relative error $\epsilon = e/\hat{w}$ is more sufficient to compare the performance in visual body weight estimation for a large variety in ground truth body weight. Additionally, the mean average error $e_{mae} = 1/n \sum_{i=1}^n |\hat{w}_i - \tilde{w}_i|$ and the mean square error $e_{mse} = 1/n \sum_{i=1}^n (\hat{w}_i - \tilde{w}_i)^2$ are evaluated.

For experiments the neural network was trained with a subset of 25 percent of the total datasets $\mathbf{f}_i \in \mathbf{F}_T \subset \mathbf{F}$ with $i = \{1, 2, \dots, 0.25 \cdot I\}$ to prevent overfitting. For validation the complete dataset was applied $\mathbf{F}^V = \mathbf{F}$. All features are rescaled in a range of $[0, 1]$. The ANN \mathbf{N} is designed as a three layer feed forward network, including one input, one hidden and one output layer. The number of input units is given by the number of features M forwarded to the ANN, while the output layer consists of a single unit for the body weight. The number of hidden units is set to the same number of input units M . A higher number of hidden units could likely slip easier into overfitting. The sigmoid function $g(x) = 1/(1+e^{-x})$ was chosen to be applied as activation function for all neurons.

The dataset for training is randomized after each trial in training which causes slower training but achieves better results. Training the network is aborted when the error in the testing dataset starts to increase. Furthermore, the training is aborted at the latest of 250 iterations. Learning is achieved by resilient propagation (Riedmiller and Braun, 1993). Regularization is applied with weight decay to improve the outcome. After training, a forwarded feature vector \mathbf{f}_i will result in an estimation for body weight \tilde{w} with

$$\tilde{w}_i = \mathbf{N}(\mathbf{f}_i) \quad . \quad (2)$$

Due to random starting points, the training was repeated with $K = 100$ trials for each experiment. For most purposes in training of neural networks the mean square error over the absolute body weight might be applied, giving a solution without any further constrains. Here, the mean

square error function is applied over the relative error based on the total dataset of features \mathbf{F} . After 100 trials the best solution \mathcal{E} is taken for experiments:

$$\mathcal{E}(\mathbf{N}_k, \mathbf{F}) = \arg \min_{\epsilon_i \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^N \epsilon_i^2 \quad (3)$$

$$\text{with } \epsilon_i = \frac{\hat{w}_i - \mathbf{N}_k(\mathbf{f}_i)}{\hat{w}_i} \quad \mathbf{f}_i \in \mathbf{F}$$

$$\text{and } k = \{1, 2, \dots, K\}$$

6. RESULTS AND DISCUSSION

The different configurations for experiments are presented and discussed in this section. Table 3 demonstrates the results from experiments. Figure 5 shows the results of all here presented experiments in a cumulative error plot.

Experiment e_1 : Volume and body weight have the highest correlation with a value of 0.93. This was the motivation for Pfitzner et al. (2015) to estimate the body weight by a volumetric reconstruction of the person in front of the sensor. To compete against this approach the volume was considered to be the only forwarded feature to the neural network. In previous work weight estimation was in a range of $\pm 10\%$ with a value of 79%. Here the results show an increase for these estimations towards 83.7%; also the standard deviation is decreased from 8.6 kg to 7.6 kg.

Experiment e_2 : Adding the frontal surface of the person also improves the approach: The range for outliers is decreased, as well as the standard deviation. Only the estimation within a range of $\pm 10\%$ is slightly decreased.

Experiment e_3 : To compare the results from Pfitzner et al. (2016) the same feature sets are forwarded to the ANN. The results are similar: The standard deviation slightly decreased from 6.5 kg to 5.8 kg in this setting. The weight estimation within a range of $\pm 10\%$ is also improved slightly from 89.8% to 90.6%. The differences might be explained due to the bigger dataset in training.

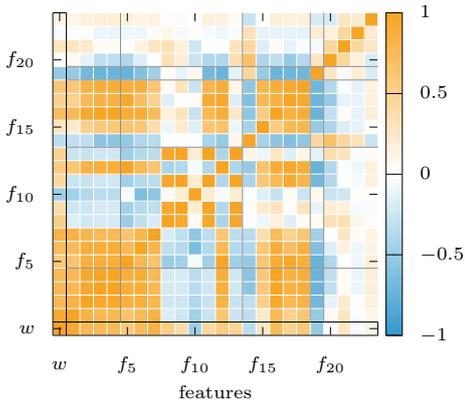


Fig. 4. Correlation of different features and ground truth \hat{w} body weight.

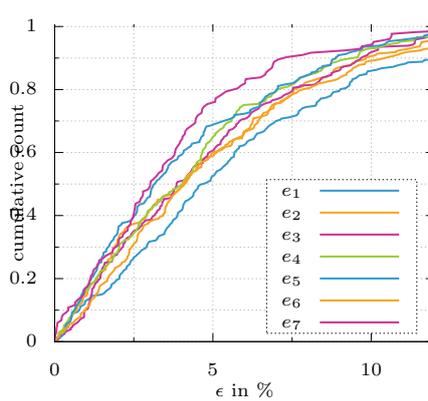


Fig. 5. Cumulative count from relative error.

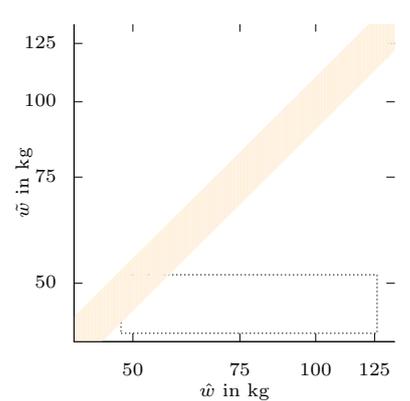


Fig. 6. Ground truth weight \bar{w} over estimated weight \hat{w} from experiment e_6 with $\pm 10\%$ margin.

Experiment e_4 : Other related work already noticed a difference in visual body weight estimation depending on the gender (Nguyen et al., 2014). Looking at the results, this approach is improved compared to experiment e_4 : The maximum outlier is reduced to 14.3kg. Furthermore, the cumulative count for $\pm 10\%$ increases from 91.0% to 91.4%. Having a look on the standard deviation σ , the mean absolute error e_{mae} and the mean squared error e_{mse} are close to the previous presented experiment and slightly improved.

Experiment e_5 : For this test all available features are used to compare the previously presented experiments. To improve the results in body weight estimation the data from the thermal imaging camera is used. Depending on the clothes of a patient the extracted features differ. Thick clothes result in a higher volume and surface. To minimize this effect the minimum, maximum and mean temperature of the patient as well as the ambient temperature are extracted and forwarded to the neural network. These additional features show only minimal correlation (see Fig. 4). Figure 6 demonstrates this result as a scatter plot.

Experiment e_6 : The volume estimation described for experiment e_1 is only possible for a laying person. To obtain the volume with a single RGB-D sensor, a standing person has to be seen from different poses, e.g. rotating around itself. Via an Iterative Closest Point (ICP) algorithm the frames from the RGB-D camera can be registered and a volumetric reconstruction can be applied Besl and McKay (1992). This approach has several disadvantages: The data acquisition takes several seconds, as well as the reconstruction. The test person might move during the data acquisition which could cause problems during registration or the ICP has to be extended to register non-rigid models Hahnel et al. (2003). Table 3 demonstrates the results for estimating a persons body weight without knowing the volume: This approach has the lowest result for the margin of 20% body weight estimation with 98.7%. Looking for the standard deviation and the range in relative error still a better estimation of body weight than seen in experiments e_1 and e_2 .

Experiment e_7 : To demonstrate the robustness of the presented features a second sensor is used for comparison. The Microsoft Kinect works with the structured light principle, having a resolution in depth of 320×240 . Compared to that the second edition of the Microsoft Kinect One works with the time of flight sensor principle and provides a resolution of 512×424 in depth image. For this experiment all available features $f_1 - f_{23}$ are used for training. Only the *Event* dataset provides data from Kinect One with 126 datasets. Therefore this experiment has to be marked as preliminary and will be verified in an upcoming publication with a higher amount of datasets. The result in body weight estimation with the Kinect One improved further: This experiment had the best minimum relative error. Further, it has the best results for the relative error in a range of $\pm 5\%$ as well as in $\pm 10\%$, see Tab.3.

7. CONCLUSION AND FUTURE WORK

This paper demonstrated the extraction of features for body weight estimation, setting a focus on robust features having a correlation to human body weight. The features rely on geometric features from eigenvalues. Experiments showed that adding the gender could improve the outcome of the body weight estimation, as well as the addition of thermal features from the person to compensate the thermal isolation in volume estimation. The best result was achieved if all available features are used for estimation and a cumulative count of 94.8% is reached for $\pm 10\%$ when experiments were done with Kinect camera. Using the Kinect One the result is improved to 95.3% for the range of $\pm 10\%$ relative error. In an up-coming clinical trial the results of this algorithm will be verified. The study will start at the Universitätsklinikum Erlangen, Germany, concerning body weight estimation for stroke patients.

For future work a dataset with a high variety in test person will be produced. Several different poses should be tracked, like standing, walking or sitting. This dataset will be made public, containing the raw RGB-D values including thermal data to encourage this area of research. Adding the skeleton model to estimate someone's pose might also improve the results further, while a bigger

Table 3. List of features used for experiments and statistical results. The best results for each column are marked in bold. For comparison the results of Pfitzner et al. (2016) and Nguyen et al. (2014) are added to this table. All listed experiments rely on data from Kinect camera, except experiment e_7 which is based on data from Kinect One.

exp.	features	samples	sensor	rel. error in %					in range in %			error in kg/kg ²	
				min	max	range	avg	σ	in 5%	in 10%	in 20%	ϵ_{mae}	ϵ_{mse}
e_1	f_1	233	Kinect	-22.9	26.8	38.0	0.21	7.6	50.6	83.7	99.1	4.68	35.3
e_2	$f_1 - f_2$	233	Kinect	-20.3	25.3	34.8	0.28	7.3	51.1	83.3	99.1	4.46	32.8
e_3	$f_1 - f_{14}$	233	Kinect	-19.0	19.8	29.7	-0.07	5.8	61.8	91.0	100.0	3.65	22.0
e_4	$f_1 - f_{16}$	233	Kinect	-14.0	14.3	26.2	0.28	5.5	63.9	91.4	100.0	3.42	19.2
e_5	$f_1 - f_{23}$	233	Kinect	-12.9	17.6	21.3	0.29	5.3	67.8	94.8	100.0	3.21	17.0
e_6	$f_2 - f_{23}$	233	Kinect	-16.8	22.7	34.2	0.43	6.6	58.4	87.1	98.7	4.00	26.9
e_7	$f_1 - f_{23}$	106	Kinect One	-8.7	14.3	23.0	0.90	4.8	75.6	95.3	100.0	2.86	13.8
Pfitzner		69	Kinect	-14.5	17.6	32.1	-0.7	6.53	–	89.9	–	–	–
Nguyen		400	Kinect	–	–	–	–	–	–	–	–	5.20	–

dataset for training is mandatory. Having a bigger dataset available can lead to a deep learning approach, without indication of a certain feature set.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the contributions from Martin Köhrmann and Lorenz Breuer from the Universitätsklinikum Erlangen, Germany, who were responsible for data acquisition.

REFERENCES

- Ailisto, H., Vildjiounaite, E., Lindholm, M., Mäkelä, S.M., and Peltola, J. (2006). Soft biometrics combining body weight and fat measurements with fingerprint biometrics. In *Pattern Recognition Letters*, volume 27, 325–334.
- Besl, P.J. and McKay, N.D. (1992). A Method for Registration of 3-D Shapes. *IEEE transactions on pattern analysis and machine intelligence*, 14, 239–256.
- Breuer, L., Nowe, T., Huttner, H.B., Blinzler, C., Kollmar, R., Schellinger, P.D., Schwab, S., and Köhrmann, M. (2010). Weight Approximation in Stroke Before Thrombolysis The WAIST-Study: A Prospective Observational Dose-Finding Study. *Stroke*, 41(12), 2867–2871.
- Buckley, R.G., Stehman, C.R., Dos Santos, F.L., Riffenburgh, R.H., Swenson, A., Mjos, N., Brewer, M., and Mulligan, S. (2012). Bedside Method to Estimate Actual Body Weight in the Emergency Department. 42(1), 100–104.
- Dantcheva, A., Velardo, C., D’angelo, A., and Dugelay, J.L. (2010). Bag of soft biometrics for person identification: New trends and challenges. *Multimedia Tools and Applications*, Springer.
- Fernandes, C.M., Clark, S., Price, A., and Innes, G. (1999). How accurately do we estimate patients’ weight in emergency departments? *Canadian Family Physician*, 45, 2373–2376.
- Fischler, M.A. and Bolles, R.C. (1981). Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Commun. ACM*, 24(6), 381–395.
- Gabel, M., Gilad-Bachrach, R., Renshaw, E., and Schuster, A. (2012). Full body gait analysis with Kinect. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 1964–1967*.
- Hahnel, D., Thrun, S., and Burgard, W. (2003). An Extension of the ICP Algorithm for Modeling Nonrigid Objects with Mobile Robots. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI’03*, 915–920. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Linder, T., Wehner, S., and Arras, K.O. (2015). Real-time full-body human gender recognition in (RGB)-D data. In *IEEE International Conference on Robotics and Automation, ICRA 2015, Seattle, WA, USA, 26-30 May, 2015*, 3039–3045. IEEE.
- Nguyen, T.V., Feng, J., and Yan, S. (2014). Seeing Human Weight from a Single RGB-D Image. *Journal of Computer Science and Technology*, 29(5), 777–784.
- Pfitzner, C., May, S., Merkl, C., Breuer, L., Braun, J., and Dirauf, F. (2015). Libra3D: Body Weight Estimation for Emergency Patients in Clinical Environments with a 3D Structured Light Sensor. *IEEE International Conference on Robotics and Automation 2015*.
- Pfitzner, C., May, S., and Nüchter, A. (2016). Neural network-based visual body weight estimation for drug dosage finding. In *Proceedings of the SPIE Medical Imaging 2016*, volume 9784, 97841Z–97841Z–9.
- Pirker, K., Bischof, H., and Skrabal, F. (2010). Human Body Volume Estimation in a Clinical Environment.
- Riedmiller, M. and Braun, H. (1993). A Direct Adaptive Method for Faster Backpropagation Learning: The RPROP Algorithm. In *IEEE International Conference on Neural Networks*, 586–591.
- Robinson, M. and Parkinson, M.B. (2013). Estimating anthropometry with microsoft kinect. *Proceedings of the 2nd International Digital Human Modeling Symposium*.
- Rusu, R.B. and Cousins, S. (2011). 3D is here: Point Cloud Library (PCL). *Proceedings - IEEE International Conference on Robotics and Automation*, 1–4.
- Shotton, J., Fitzgibbon, A., Cook, M., Sharp, T., Finocchio, M., Moore, R., Kipman, A., and Blake, A. (2013). Real-time human pose recognition in parts from single depth images. *Studies in Computational Intelligence*, 411, 119–135.
- Velardo, C. and Dugelay, J.L. (2012). What can computer vision tell you about your weight? In *EUSIPCO 2012, 20th European Signal Processing Conference 2011, August 27-31, 2012, Bucharest, Romania*. Bucharest, Romania.