Scan-over-Clothes (SOC): Improved Body Measurement Accuracy When Scanning Loose-Clothed Subjects

Matthew S. GILMER, Steven C. HAUSER David BRUNER
Size Stream LLC, Cary NC, USA.

https://doi.org/10.15221/25.04

Abstract

At Size Stream we continuously strive to enhance the accuracy of our 3D body scanning technology while minimizing user friction. We have previously employed a model that detects loose clothing during scanning. In the cases of positive detection, we have prompted the user with a suggestion to change into tighter-fitting attire. We now introduce a method that effectively compensates for loose clothes for our avatar generation and measurement predictions. For scans of subjects wearing clothing that is marginally loose and only partially covering the body, this compensation results in an accuracy degradation of less than 10% compared to scans of those same subjects wearing ideal scanning attire. This improvement was achieved through two primary developments. First, we developed a `Body-Plus-Clothes (BPC)` segmenter which, in addition to separating the subject from the background, distinguishes the clothing from the bare body. Aside from the added capability to segment clothing, this model also provides a substantial accuracy improvement in separating the subject from the background when compared to our previous segmenter. The `BPC` segmenter was trained using real images and supervised to manually-corrected output from an open-source segmentation model. Second, leveraging the BPC segmenter, we developed a `Scan-Over-Clothes (SOC)` Body Measurement Model (BMM) specifically designed to adjust for the presence of loose clothing during body reconstruction and measurement estimation. The `SOC` model was trained using a combination of `BPC` segmentation of real images and internally-generated synthetic data. This novel approach results in a solution that substantially enhances measurement accuracy for scans involving loose clothing, raising the possibility for reliable Ready-To-Wear (RTW) sizing with relaxed scan wear requirements.

Keywords: 3D body scanning, machine learning, mobile scanning, segmentation, synthetic data

1. Introduction

We at Size Stream have long been providing Made-To-Measure (MTM) clothing to our customers. When we started, we required measurements from one of our booth scanners. The final version of this was the `SS20`, which followed established 3D body measurement practices [1, 2]. Beginning in 2020, we shifted our focus to mobile measurement solutions to make the scanning process more accessible. After years of improving the measurement quality in our mobile scanning, we achieved a 7% remake rate for MTM clothing. Our mobile scanning technology was evaluated alongside the `SS20` with each showing strong correspondence in anthropometric circumferences to manual tape measurements [3]. Accurate and repeatable anthropometric measurements are critical for MTM apparel systems, where even small systematic errors can lead to fit dissatisfaction and garment remakes [4].

We are now exploring another avenue towards making the scan process easier. In this paper we describe our first steps towards relaxing our scan wear requirements. We introduce two key innovations: the `Body-Plus-Clothes (BPC)` segmenter, capable of distinguishing a subject's body from their worn clothing with high fidelity, and the `Scan-Over-Clothes (SOC)` Body Measurement Model (BMM). The latter is trained to compensate for loose garments using BPC segmentation as input, extending the application space relative to that of prior machine-learning techniques for measurement regression from segmentation, which typically assume silhouettes minimally distorted by clothing and are therefore sensitive to loose clothing [5]. As a bonus, we find that the `BPC` segmenter significantly outperforms our previous segmenter (`S2023`) in the task of separating the subject from the background. While we have yet to put the `SOC` BMM into practice, we began upgrading our mobile apps to use the `BPC` segmenter for subject segmentation in June of this year.

This paper is organized as follows: in Section 2, we describe the methods by which we developed these key innovations beginning with the `BPC` segmenter (Section 2.1) and ending with the `SOC` BMM (Section 2.2). In Section 3 we present the results following the same structure as Section 2. Finally, we summarize our findings in Section 4.

2. Methods

2.1. Body-Plus-Clothes (BPC) Segmenter

Here we describe the development of the `BPC` segmenter that we first deployed in June. It is currently only being used for the task of separating the subject from the background, but it also separately segments clothing worn by the subject. We achieved this behavior by training on data in which the ground truth consisted of pixel-level labels corresponding to the bare body, worn clothing, and the background. The training data were drawn from two sources: (1) a publicly available fashion dataset comprising 45000 pre-annotated images, and (2) 10515 images collected through our body-scanning applications in accordance with Size Stream privacy agreements. The latter did not initially have the required ground-truth labels, and we describe the process of obtaining these in Section 2.1.1.

2.1.1. Preparation of Training Data

The first step in obtaining the ground truth for our body-scanning images was to batch process the images using an open-source segmentation engine, which we cannot name due to corporate policy. The raw batch process output was initially not suitable for training as articles of clothing were often missed (partially or entirely) and/or background objects would be segmented as clothing, therefore results needed to be reviewed and, when needed, corrected.

We had previously developed a Graphical User Interface (GUI) for data annotation which we used to prepare training data for a pose detector and our real-time scan quality checking system (`CHECKER`) [6], it will hereafter be referred to as `Annotation GUI (AGUI)` and is depicted in Figure 1. We adapted `AGUI` to this new task in the following ways. First, we configured `AGUI` to display the RGB image with a single segmentation result overlaid (red for clothes or blue for body). The annotator can cycle through segmentation results for each image, choosing to discard them or not using the top panel of checkboxes, and then finally to the next image using keystrokes. Second, we added functionality for the annotator to draw a line over the image by clicking and dragging the mouse. If the annotator made a mistake, they could use the keyboard to clear the entire line or clear the last drawn point. The line drawn specifies a new blob in which the final point connects in a straight line to the initial point. The annotator could use keystrokes and this new blob to either add to or subtract from the displayed segmentation blob, with the display updating immediately.

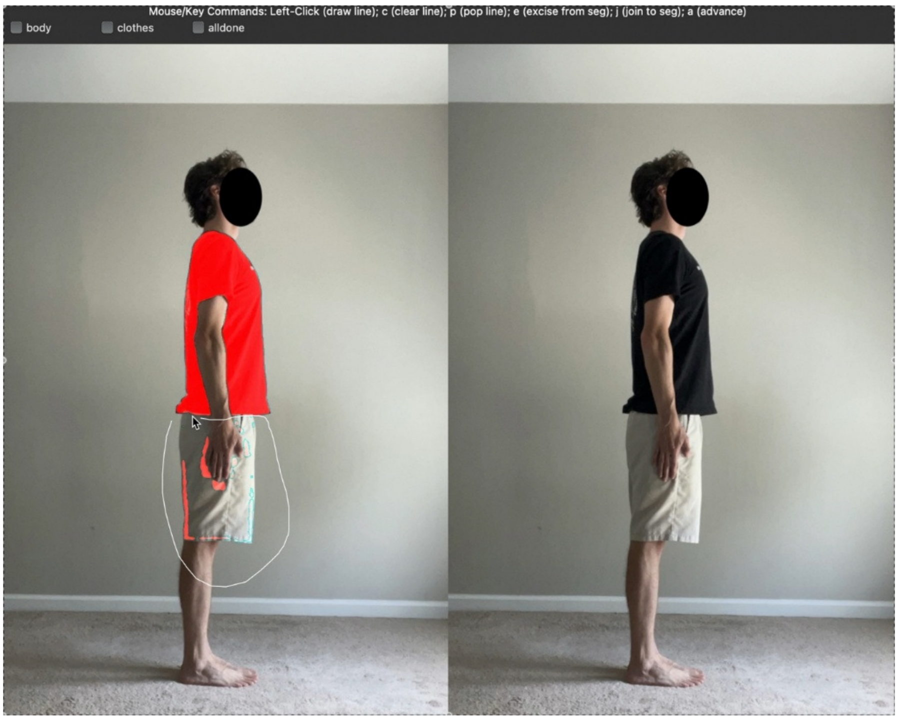

Fig. 1. Screenshot from `AGUI`. The left side shows a segmentation result from the batch process overlaid on an image. There was some unwanted segmentation in the shorts which the reviewer was about to excise using the key 'e', see 'Mouse/Key Commands' in the top panel. See also the checkboxes in the top panel for labelling segmentation results. The right shows the same image with no segmentation result overlaid.

2.1.2. Model Training

We chose a lightweight architecture suitable to the task and to mobile deployment. We trained the model in two steps. First, we trained the model using the aforementioned publicly available fashion dataset, which we cannot name due to corporate policy. Then, we fine-tuned the model on our body-scanning dataset using standard image augmentation techniques (mirroring, translation, rotation, scaling, and color jitter), taking care to preserve the precise overlap of the body and clothing masks by performing the same augmentations on each (except in the case of color jitter).

2.2. Scan-Over-Clothes (SOC) Body Measurement Model (BMM)

The task here was to build a model capable of regressing measurements, using `BPC` segmentation as input, that will give accurate results for subjects wearing any type of clothing. Ideally, it would achieve the same result on a particular subject whether they are wearing tight-fitting or loose-fitting clothing during the scan.

To this end we used our state-of-the-art BMM, `BMM2024` (not published), to generate ground-truth measurements on our training scans. This model was distilled from our previous BMM, `BMM2024`, onto a lighter weight architecture for the purpose of reducing processing time and it takes as input the front and side silhouettes along with the height, age, weight, and gender of the subject. Here, we used a similar strategy by training with the `BMM2024` measurement results as ground-truth but now leveraging the newly developed `BPC` segmenter. We modified the architecture to allow for the input of clothing masks in addition to subject silhouettes so that the model could distinguish between parts of the silhouettes that are clothed and parts of the silhouettes that are not.

To ensure we trained to accurate measurements for the subjects in our training scans, we trained only on ‘clean’ scans, in which the subjects had worn tight-fitting clothing. Scans in which the subjects wore loose clothing were labelled ‘loose’ and only the clothing masks were used. More details on this training strategy are given in Sections 2.2.1 and 2.2.2.

2.2.1. Preparation of Training Data

We processed 11743 RGB images collected from our body-scanning applications with the `BPC` segmenter. These images were from body scans where for each scan we had two images (a front view and a side view). However, we wanted to train only with ‘clean’ scans (in which subjects were wearing tight-fitting clothing) for which the `BMM2024` results accurately reflected the subject’s body measurements. To determine which scans were ‘clean’ and which were ‘loose’ instead, we used our previously developed loose clothes detector (one part of our real-time scan quality checking scheme, `CHECKER`) [6]. While we did not train with the ‘loose’ scans, we did use the clothing masks from their BPC segmentation results. Our data augmentation scheme in this case relies on randomly sampling from the clothing masks of the loose-clothed scans to be input along with ‘clean’ silhouettes.

2.2.2. Model Training

During training, the model always receives two silhouettes (one for each of the front and side views) derived from a single ‘clean’ scan. The corresponding clothing masks for the training example are provided in one of two ways: in some cases, the clothing masks are taken from the same clean scan; in others, they are drawn from a randomly selected ‘loose’ scan. This sampling strategy ensures the model will learn to predict accurate measurements both in the presence and absence of loose clothing.

Then, we apply standard image augmentation techniques as described in Section 2.1.2 (but without color jitter). Following image augmentation, the silhouettes are explicitly modified so that their foreground pixels include all the same pixels as the accompanying clothing masks. This is done so that the training reflects the conditions under which the fully trained model will be used in practice. This training strategy generates a continuously varying stream of synthetic examples that span both clean and loose-clothed scanning conditions.

3. Results

3.1. Body-Plus-Clothes (BPC) Segmenter

We found our new segmenter exceeded our expectations in being able to reliably segment out clothing from the body. Additionally, it showed a marked improvement over `S2023` in subject segmentation (separating the scanning subject from the background). Consequently, we are in the process of upgrading our scanning applications with the `BPC` segmenter, giving our measurement results a boost in both accuracy and precision. An example scan and its segmentation result is shown in Figure 2.

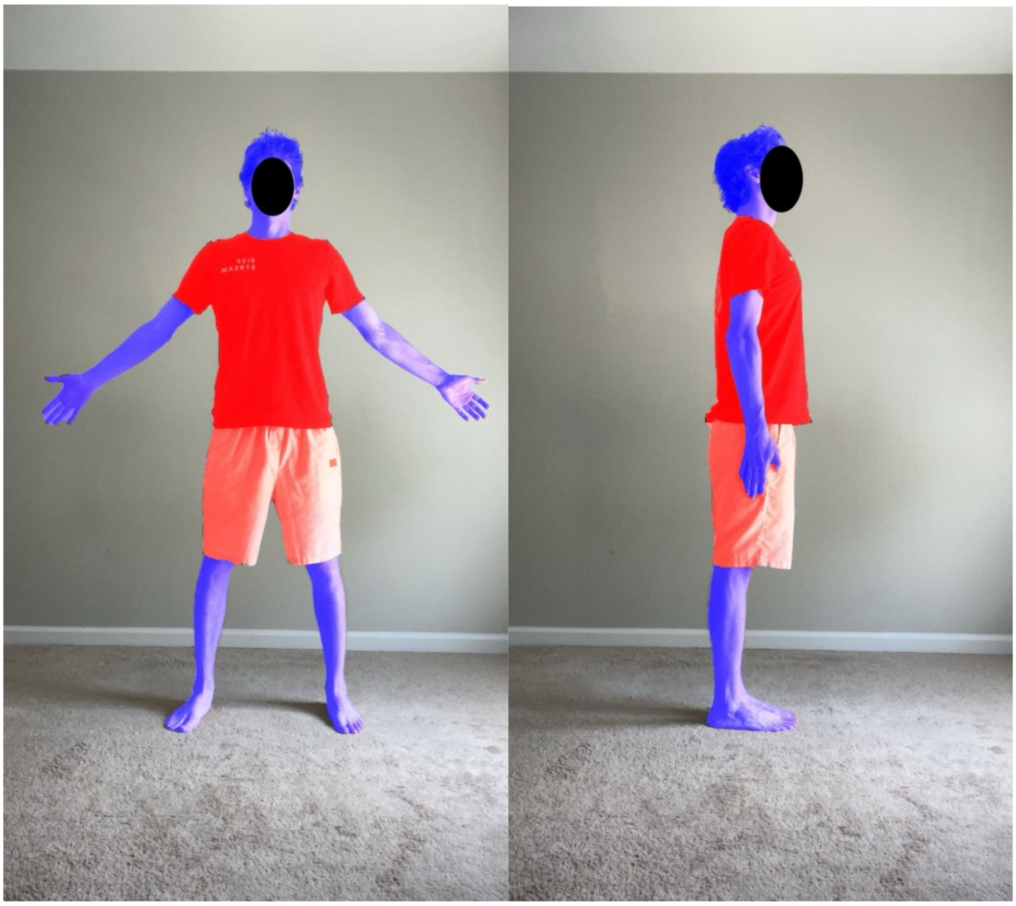

Fig. 2. Example scan with `BPC` segmentation masks represented (red: clothes, blue: body).

3.1.1. Accuracy of Subject Segmentation

The accuracy improvement in subject segmentation is reflected in the results from one of our internal 'scan party' datasets (hereafter referred to as `Dataset 1`) containing 5 mobile scans and 2 `SS20` scans for each of the 30 male and 30 female subjects (see Table 1). The number of scans and subjects were chosen to satisfy the requirement of statistical significance. The mobile scans were self-administered in a clutter-free environment at the Size Stream office, and the subjects wore skin-tight clothing for all scans. The measurements for the mobile scans were calculated using `BMM2025`. Table 1 shows the Coefficient of Variation (CV) which is the standard deviation of measurements across the 5 mobile scans of the same subject divided by their mean and expressed as a percentage. Also shown in Table 1 is the dimensionless standard deviation (dimSTD) of percent errors calculated relative to the mean of the `SS20` measurement. The CV characterizes the variability while dimSTD characterizes the accuracy, and both are averaged over all subjects and a set of 46 measurements. Additionally, we include in parentheses the percent change for the `BPC` numbers against those of the `S2023`. We find the largest improvement between `S2023` and the `BPC` segmenter is in the CV, but the data also show an accuracy improvement of 2 percent.

Table 1. Variability (CV) and accuracy (dimSTD) comparison between the segmenters `S2023` and `BPC` on `Dataset 1`.

Gender	Segmenter	CV [%]	dimSTD [%]
Male	S2023	0.68	4.13
	BPC	0.60 (-12%)	4.05 (-2%)
Female	S2023	0.63	3.56
	BPC	0.54 (-14%)	3.48 (-2%)

The positive impact on our customers who scan at home is more difficult to quantify since we do not have `SS20` scans for them and scans are not taken in immediate succession by single subjects, however, it is likely stronger due to the prevalence of lighting issues and cluttered scan environments (absent from `Dataset 1`), which the `BPC` segmenter has improved resiliency against (see Figure 3). In lieu of a quantitative evaluation we inspected qualitatively a random sample of 100 scans (200 images) using both the `S2023` and `BPC` segmenters. We found only two examples of `BPC` segmentation where there was a small portion of the background segmented as part of the subject and in each case the real boundary was difficult to see by eye due to background clutter. Additionally, we noted that the `BPC` segmenter perfectly segmented limbs with tattoos that `S2023` could not and it routinely produces segmentations with sharper silhouette boundaries and less background noise.

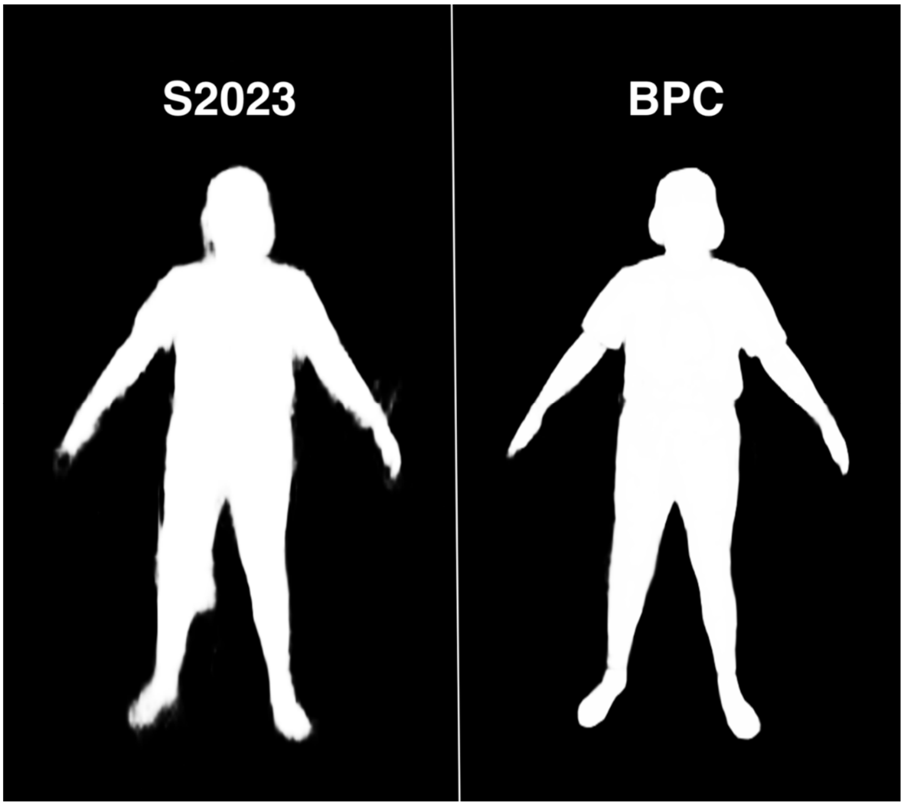

Fig. 3. Comparison of the `S2023` (left) and the `BPC` (right) segmentation results for an example front view image that had bad lighting and a cluttered background.

3.1.2. Accuracy of Clothing Segmentation

In the same sample of 100 scans (see Section 3.1.1) we found only two examples where a garment was partially missed, in both cases the missed region was counted as body but not clothes. The model also reliably segmented relatively rare garments like hats. We concluded that the clothing segmentation of the `BPC` segmenter is reliable enough to be used as input for a SOC BMM.

3.2. Scan-Over-Clothes (SOC) Body Measurement Model (BMM)

We evaluate our `SOC` BMM on two different datasets. The first is `Dataset 1` (introduced in Section 3.1.1). This dataset has only 'clean' scans, but we include the evaluation here since we trained the model to be able to handle both 'clean' and 'loose' scans. The second is `Dataset 2` which contains 10 mobile scans and 2 `SS20` scans for each of the 8 male and 7 female subjects. The number of subjects in this dataset was limited by resources; however, the results are strong enough to make the analysis significant. Half of the 10 mobile scans were 'clean' and half were 'loose'. For the 'loose' scans, the subjects were allowed to choose their clothing which resulted in a mix of shorts, slacks, jeans, sweatpants, t-shirts, and sweaters.

3.2.1. Evaluation on Dataset 1

The results for this evaluation are shown in Table 2. We also include the percent changes of all metrics for the `SOC` BMM against those of `BMM2025` (negative in green and positive in red). We find that for the `SOC` BMM the CV is slightly lower for the male subjects and slightly higher for the female subjects. We also find that the `SOC` model has dimSTD values that are slightly higher (with percent changes not exceeding 10%) for both male and female genders. Therefore, the `SOC` BMM is slightly less accurate on 'clean' scans than BMM2025.

Table 2. Variability (CV) and accuracy (dimSTD) comparison between the BMMs `BMM2025` and `SOC` on `Dataset 1`.

Gender	BMM	CV [%]	dimSTD [%]
Male	BMM2025	0.60	4.05
	SOC	0.57 (-5%)	4.34 (+7%)
Female	BMM2025	0.54	3.48
	SOC	0.57 (+6%)	3.83 (+10%)

3.2.2. Evaluation on Dataset 2

The results for this evaluation are shown in Table 3. Here the percent changes shown are with respect to the 'clean' scans run through `BMM2025`. For both male and female genders, the `SOC` BMM is only slightly less accurate than `BMM2025` on 'clean' scans. We also note the CV for female subjects on 'clean' scans is significantly lower than it is for `BMM2025`. These results, however, have less significance than the analogous results of 3.2.1 due to `Dataset 2` having a quarter the number of subjects as `Dataset 1`.

When processing scans with `BMM2025` we saw a strong decrease in accuracy for the 'loose' scans relative to the 'clean' scans of the same subjects due to `BMM2025` not being equipped to account for loose clothes in its measurement predictions. We also note that this accuracy degradation is stronger for the female subjects, and this is likely due to these subjects choosing looser clothing on average than the male subjects.

Finally, we describe the main result which is that when the `SOC` BMM is used on 'loose' scans instead of `BMM2025` all metrics are significantly reduced. The `SOC` BMM results still show a degradation in performance relative to the `BMM2025` results on the 'clean' scans (excepting the reduction in CV for male subjects) but this degradation is dwarfed by that of the results for using `BMM2025` on those same 'loose' scans. As a reminder the metrics in these tables have been averaged over a set of 46 measurements. The measurements averaged vary in their sensitivity to the 'loose' clothing, so we additionally include Figure 4 to highlight three circumference measurements that have a strong sensitivity. The biases, shown here as dimensionless mean differences (dimMD) introduced by 'loose' clothing are effectively eliminated when the `SOC` BMM is used. While we have not yet attempted to do RTW sizing with the `SOC` BMM, we are encouraged by the results presented here to expect it would prove successful. In the case of the male subjects, the accuracy degrades only 10% relative to using `BMM2025` on 'clean' scans (a method already proven to work for MTM clothing).

Table 3. Variability (CV) and accuracy (dimSTD) comparison between the BMMs `BMM2025` and `SOC` for 'clean' and 'loose' scans on `Dataset 2`.

Gender	'Clean' or 'Loose'	BMM	CV [%]	dimSTD [%]
Male	'Clean'	BMM2025	0.49	3.62
		SOC	0.49 (+0%)	3.65 (+1%)
	'Loose'	BMM2025	0.58 (+18%)	5.20 (+44%)
		SOC	0.38 (-22%)	3.97 (+10%)
Female	'Clean'	BMM2025	0.52	3.82
		SOC	0.41 (-21%)	3.90 (+2%)
	'Loose'	BMM2025	0.78 (+51%)	5.66 (+48%)
		SOC	0.56 (+8%)	4.40 (+15%)

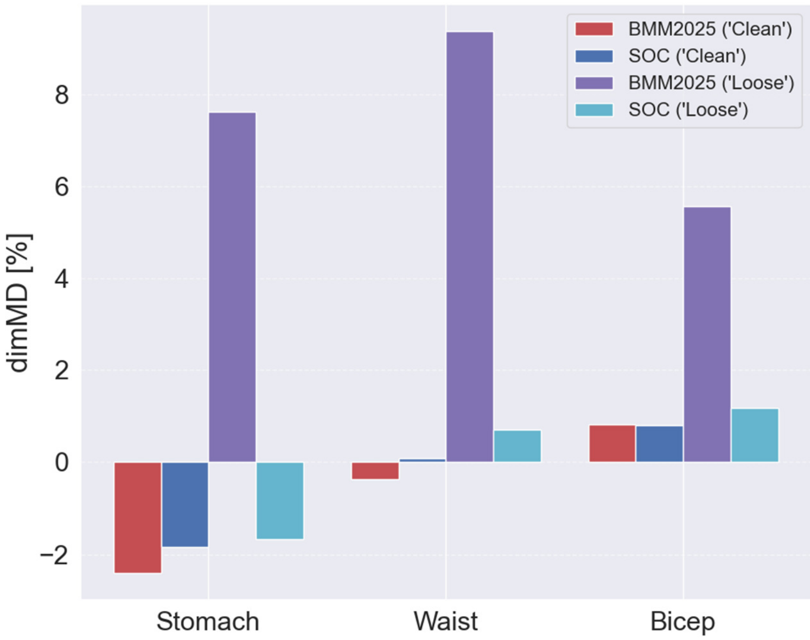

Fig. 4. dimMD (bias) of select circumferences (Stomach, Waist, and Bicep) and their associated dimensionless metrics BMMs `BMM2025` and `SOC` for 'clean' and 'loose' scans on `Dataset 2`.

4. Summary

We introduced a clothing-aware BMM (`SOC`) that greatly enhances accuracy compared to our current BMM (`BMM2025`) when subjects are scanned in loose clothing. The `SOC` BMM can additionally be used when subjects are wearing proper (skin-tight) clothing with only a slight loss of accuracy (relative to `BMM2025`). This result was made possible by the development of the `BPC` segmenter whose ability to distinguish between a subject's body and their clothing provided the information necessary for the `SOC` BMM to learn to compensate for loose clothing.

As an additional bonus to the effort in developing the `BPC` segmenter, we have significantly improved our subject segmentation (previously handled by `S2023`). The subject segmentation with the `BPC` segmenter is more robust to cluttered scan environments and poor lighting. The `BPC` segmenter also, when used with our current BMM (`BMM2025`), provides a significant improvement in precision and a modest improvement in accuracy to our reported measurements in ideal scanning conditions.

Future work may include an RTW sizing test. If we find we can provide reliable RTW sizing with relaxed clothing restrictions, which appears likely, we may explore further applications in the apparel and health spaces.

References

[1] K. M. Robinette, S. Blackwell, H. Daanen, M. Boehmer, S. Fleming, T. Brill, and D. Burnsides, "Civilian American and European Surface Anthropometry Resource (CAESAR) Final Report, Volume I: Summary", Wright-Patterson AFB, OH, USA: Air Force Research Laboratory, 2002.

[2] J. L. Istook and H. D. Hwang, "3D body scanning systems with application to the apparel industry," in *Journal of Fashion Marketing and Management: An International Journal,* Vol. 5, No. 2, 2001, pp.120-132,

[3] B. Smith, C. McCarthy, M. E. Dechenaud, M. C. Wong, J. Shepherd and S. B. Heymsfield, "Anthropometric evaluation of a 3D scanning mobile application", *Obesity,* Silver Spring, MD, Vol. 30, No. 6, Jun. 2022, pp.1181–1188, https://doi.org/10.1002/oby.23434.

[4] S. P. Ashdown, "Sizing in Clothing: Developing Effective Sizing Systems for Ready-to-Wear and Custom Clothing", 2nd ed. Hoboken, NJ, USA: Wiley, 2007.

[5] N. Ruiz, M. Bellver, T. Bolkart, A. Arora, M. C. Lin, J. Romero, and R. Bala, "Human Body Measurement Estimation with Adversarial Augmentation," *arXiv preprint*, 2022, https://arxiv.org/abs/2210.05667.

[6] M. S. Gilmer, S. C. Hauser, and D. Bruner, "CHECKER: Real-time scan quality checking and associated user experience," in *Proc. 3DBODY.TECH 2023 - 14th Int. Conf. and Exh. on 3D Body Scanning and Processing Technologies*, Lugano, Switzerland, 2023, https://doi.org/10.15221/23.30.