

Skin Lesion Correspondence Localization in Total Body Photography

Wei-Lun Huang^{1*}, Davood Tashayyod⁵, Jun Kang³, Amir Gandjbakhche⁴,
Michael Kazhdan^{1**}, and Mehran Armand^{1,2**}

¹ Johns Hopkins University, Department of Computer Science, Baltimore, MD, USA

² Johns Hopkins University, Department of Orthopaedic Surgery, Baltimore, MD,
USA

³ Johns Hopkins School of Medicine, Department of Dermatology, Baltimore, MD,
USA

⁴ Eunice Kennedy Shriver National Institute of Child Health and Human
Development, Bethesda, MD, USA

⁵ Lumo Imaging, Rockville, MD, USA

Abstract. Longitudinal tracking of skin lesions - finding correspondence, changes in morphology, and texture - is beneficial to the early detection of melanoma. However, it has not been well investigated in the context of full-body imaging. We propose a novel framework combining geometric and texture information to localize skin lesion correspondence from a source scan to a target scan in total body photography (TBP). Body landmarks or sparse correspondence are first created on the source and target 3D textured meshes. Every vertex on each of the meshes is then mapped to a feature vector characterizing the geodesic distances to the landmarks on that mesh. Then, for each lesion of interest (LOI) on the source, its corresponding location on the target is first coarsely estimated using the geometric information encoded in the feature vectors and then refined using the texture information. We evaluated the framework quantitatively on both a public and a private dataset, for which our success rates (at 10 mm criterion) are comparable to the only reported longitudinal study. As full-body 3D capture becomes more prevalent and has higher quality, we expect the proposed method to constitute a valuable step in the longitudinal tracking of skin lesions.

Keywords: Total body photography · Skin lesion longitudinal tracking · 3D correspondence.

1 Introduction

Evolution, the change of pigmented skin lesions, is a risk factor for melanoma [1]. Therefore, longitudinal tracking of skin lesions over the whole body is beneficial for early detection of melanoma [5]. However, establishing skin lesion

* Corresponding author: Wei-Lun Huang (whuang44@jh.edu)

** Co-senior authors

correspondences across multiple scans from different patient visits has not been well investigated in the context of full-body imaging.

Several techniques have been proposed to match skin lesions across pairs of 2D images [9,10,11,12,13,14,16,17,25]. Early work used geometric constraints imposed by initial matches of skin lesions (manual selection or automatic detection) to align images and further match other skin lesions [10,16,17,25]. Mirzaalian and colleagues published a series of works for establishing lesion correspondence in image space [11,12,13,14]. Li et al. [9] used a CNN to output a 2D vector field for pixel-wise correspondences between the two input images. Though effective at matching skin lesions across pairs of images, the extension of these methods to the context of total body photography (TBP) for longitudinal tracking remains a challenge.

Several works have been proposed for tackling the skin lesion tracking problem over the full body [7,8,21,22]. However, they are either only applicable in well-controlled environments or do not extend to the tracking of lesions across scans at different visits. Recently, the concept of finding lesion correspondence using a 3D representation of the human body has been explored in [26] and [2] by using a template mesh. However, accurately deforming a template mesh to fit varying body shapes is challenging when the scanned shape deviates from the template, leading to large errors in downstream tasks such as establishing shape correspondence. Additionally, [26] does not take advantage of texture, while [2] uses texture in a common UV map that may lead to failures when geodesically close locations on the surface are mapped to distant sites in the texture map (e.g. when the two locations are on opposite sides of a texture seam).

We propose a novel framework for finding skin lesion correspondence iteratively using geometric and texture information (Fig. 1). We demonstrate the effectiveness of the proposed method in localizing lesion correspondence across scans in a manner that is robust to changes in body pose and camera viewing directions. Our code is available at <https://github.com/weilunhuang-jhu/LesionCorrespondenceTBP3D>.

2 Methods

Given a set of lesions of interest (LOIs) X in the source mesh, we would like to find their corresponding positions Y in the target mesh. Formally, we assume we are given source and target meshes, \mathcal{M}_0 and \mathcal{M}_1 , with vertex sets $\mathcal{V}_k \subset \mathcal{M}_k$ and corresponding landmark sets $L_k \subset \mathcal{V}_k$ with $|L_0| = |L_1| = S$. We achieve this by computing a dense correspondence map $\Phi_{L,\epsilon} : \mathcal{V}_0 \rightarrow \mathcal{V}_1$, initially defined using geometric information and refined using textural information. Then we use that to define a map taking lesions of interest on the source to positions on the target $\Phi : X \rightarrow \mathcal{V}_1$.

2.1 Landmark-based correspondences

We define an initial dense correspondence between source and target vertices by leveraging the sparse landmark correspondences [3,6]. We do this by mapping

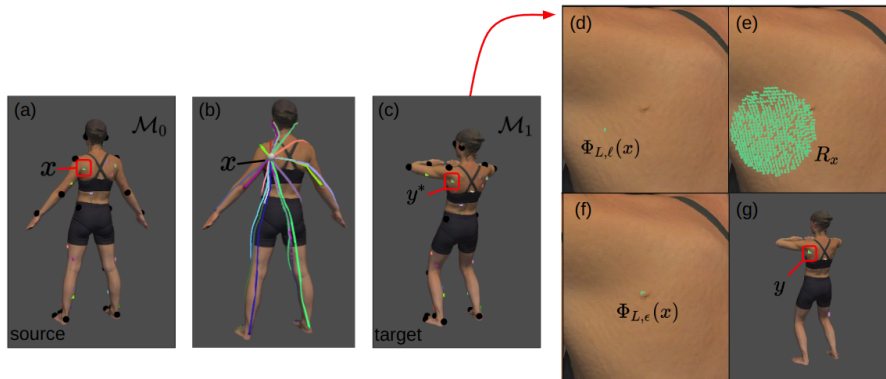


Fig. 1. Visualization of the correspondence localization workflow with geometric and texture information. (a) and (c) show a source and a target textured mesh with landmarks (black dots) and labeled lesions of interest (color dots). (b) shows the geodesic paths from a lesion x in (a) to all the landmarks. (d) shows the correspondence of the lesion $\Phi_{L,l}(x)$ derived from the geometric feature descriptors. (e) shows the local region R_x for refining the position of the correspondence. (f) shows the correspondence of the lesion $\Phi_{L,e}(x)$ from the local texture descriptors. (g) shows the lesion correspondences with confidence and anchored as new landmarks after one iteration.

source and target vertices into a high-dimensional space, based on their proximity to the landmarks, and then comparing positions in the high-dimensional space.

Concretely, we define maps $\ell_k : \mathcal{V}_k \rightarrow \mathbb{R}^S$, associating a vertex $v \in \mathcal{V}_k$ with an S -dimensional feature descriptor that describes the position of v relative to the landmarks:

$$\ell_k(v) = \left[\frac{1}{D_k(v, l_1)}, \dots, \frac{1}{D_k(v, l_S)} \right]^\top \quad (1)$$

where $D_k : \mathcal{M}_k \times \mathcal{M}_k \rightarrow \mathbb{R}^{\geq 0}$ is the geodesic distance function on \mathcal{M}_k . We use the reciprocal of geodesic distance so that landmarks closer to v contribute more significantly to the feature vector.

Given this mapping, we create an initial dense correspondence between the source and target vertices, $\Phi_{L,l} : \mathcal{V}_0 \rightarrow \mathcal{V}_1$ by mapping a source vertex $v \in \mathcal{V}_0$ to the target vertex with the most similar feature descriptor (with similarity measured in terms of the normalized cross-correlation):

$$\Phi_{L,l}(v) = \arg \max_{v' \in \mathcal{V}_1} \left\{ C_{L,l}(v, v') = \frac{\langle \ell_0(v), \ell_1(v') \rangle}{\|\ell_0(v)\| \cdot \|\ell_1(v')\|} \right\}. \quad (2)$$

2.2 Texture-based refinement

While feature descriptors of corresponding vertices on the source and target mesh are identical when 1) the landmarks are in perfect correspondence, and 2)

the source and target differ by an isometry, neither of these assumptions holds in real-world data. To address this, we use local texture to assign an additional feature descriptor to each vertex and use these texture-based descriptors to refine the coarse correspondence given by $\Phi_{L,\ell} : \mathcal{V}_0 \rightarrow \mathcal{V}_1$. Various texture descriptors have been proposed, e.g. SHOT [23,24,20], RoPS[4], and ECHO[15]. We selected the ECHO descriptor for its better descriptiveness and robustness to noise.

Letting $\epsilon_k(v) \in \mathbb{R}^N$ denote the ECHO descriptor of vertex $v \in \mathcal{V}_k$, our goal is to refine the dense correspondence so that corresponding source and target vertices also have similar descriptors. However, to avoid problems with repeating (local) textures, we would also like the correspondence to stay close to the correspondence defined by the landmarks.

We achieve this as follows: To every source vertex $v \in \mathcal{V}_0$ we associate a region $R_v \subset \mathcal{V}_1$ of target vertices that are either close to $\Phi_{L,\ell}(v)$ (the corresponding vertex on \mathcal{V}_1 as predicted by the landmarks) or have similar geometric feature descriptors:

$$R_v = \left\{ v' \in \mathcal{V}_1 \mid D_1(v', \Phi_{L,\ell}(v)) < \varepsilon_1 \text{ or } C_{L,\ell}(v', \Phi_{L,\ell}(v)) > \varepsilon_2 \right\}. \quad (3)$$

Given this region, we define the target vertex corresponding to a source as the vertex within the region that has the most similar ECHO descriptor (using the normalized cross-correlation as before).

In practice, we compute the ECHO descriptor over three different radii, obtaining three descriptors for each vertex, $\epsilon_k^1(v), \epsilon_k^2(v), \epsilon_k^3(v) \in \mathbb{R}^N$. The selection of three different radii in ECHO descriptors is done to accommodate different sizes of lesions and their surrounding texture, and the values are empirically determined. This gives a mapping $\Phi_{L,\epsilon} : \mathcal{V}_0 \rightarrow \mathcal{V}_1$ defined in terms of the weighted sum of cross-correlations:

$$\Phi_{L,\epsilon}(v) = \arg \max_{v' \in R_v} \left\{ C_{L,\epsilon}(v, v') = \sum_{i=1}^3 w_i \frac{\langle \epsilon_0^i(v), \epsilon_1^i(v') \rangle}{\|\epsilon_0^i(v)\| \cdot \|\epsilon_1^i(v')\|} \right\}, \quad (4)$$

where $C_{L,\epsilon}(v, v') \in [0, 1]$ is the texture score of the target vertex v' and w_i is the weight of the cross-correlation between the ECHO descriptors computed at each radius.

2.3 Iterative Skin Lesion Correspondence Localization Framework

While each source LOI has a corresponding position on the target mesh as given by $\Phi_{L,\epsilon} : \mathcal{V}_0 \rightarrow \mathcal{V}_1$, not all correspondences are localized with high confidence when 1) the local texture is not well-preserved across scans and 2) the local region R_v does not include the true correspondence. To address this, we adapt our algorithm for computing the correspondence map $\Phi : X \rightarrow \mathcal{V}_1$ by iteratively growing the set of landmarks to include LOI correspondences about which we are confident, similar to the way in which a human annotator would label lesion correspondence (Fig. 2).

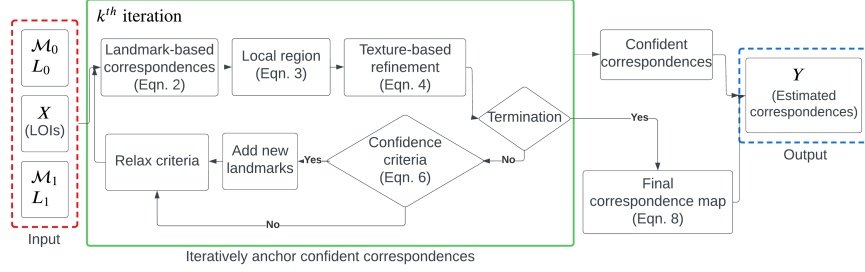


Fig. 2. Block diagram of the iterative lesion correspondence localization algorithm.

Iteratively anchor confident correspondences We iteratively compute correspondence maps $\Phi_{L,\ell}^k : \mathcal{V}_0 \rightarrow \mathcal{V}_1$ and $\Phi_{L,\epsilon}^k : \mathcal{V}_0 \rightarrow \mathcal{V}_1$, with the superscript denoting the k^{th} iteration. For each map $\Phi_{L,\epsilon}^k$ and every LOI $x \in X$, we determine if we are confident in the correspondence $\{x, \Phi_{L,\epsilon}^k(x)\}$ by evaluating a binary function $\chi_L^k : X \rightarrow \{0, 1\}$. Denoting by X' the subset of LOIs about which we are confident, we add the pairs $\{x', \Phi_{L,\epsilon}^k(x')\}$ to the landmark set L and remove the LOI $x' \in X'$ from X . We iterate this process until all the correspondences of LOIs are confidently found or a maximum number of iterations (K) have been performed.

Lesion correspondence confidence is measured using three criteria: i) texture similarity, ii) agreement between geometric and textural correspondences, and iii) the unique existence of a similar lesion within a region. To quantify uniqueness, we compute the set of target vertices whose textural descriptor is similar to that of the LOI:

$$S_x^\delta = \{v' \in R_x \mid C_{L,\epsilon}(x, v') > \delta\} , \quad (5)$$

and consider the diameter of the set (defined in terms of the mean of the distances of vertices in S_x^δ from the centroid of S_x^δ). Putting this together, we define confidence as

$$\chi_L^k(x) = C_{L,\epsilon}(x, \Phi_{L,\epsilon}^k(x)) > \varepsilon_3 \quad \vee \quad D_1(\Phi_{L,\ell}^k(x), \Phi_{L,\epsilon}^k(x)) < \varepsilon_4 \quad \vee \quad \varnothing(S_x^\delta) < \varepsilon_5, \quad (6)$$

where the initial values of thresholds ε_i are empirically chosen. To further support establishing correspondences, we relax the thresholds ε_i in subsequent iterations, allowing us to consider correspondences that are further away and about which we are less confident.

Final correspondence map Having mapped every high-confidence LOI to a corresponding target vertex, we must complete the correspondence for the remaining low-confidence LOIs. We note that for a low-confidence LOI $x \in X$,

the texture in the source mesh is not well-matched to the texture in the target, for any $v' \in R_x$. (Otherwise the first term in χ_L^k would be large.)

To address this, we would like to focus on landmark-based similarity. However, by definition of R_x , for all $v' \in R_x$, we know that the landmark descriptors of x and v' will all be similar, so that $C_{L,\ell}$ will not be discriminating. Instead, we use a standard transformation to turn distances into similarities. Specifically, we define geometric score between a source LOI x and a target vertex $v' \in R_x$ in terms of the geodesic distance between v' and the corresponding position of x in \mathcal{V}_1 , as predicted by the landmark descriptors:

$$C_L(x, v') = e^{-\frac{1}{2\sigma^2} D_1^2(v', \Phi_{L,\ell}^K(x))} \in [0, 1], \quad (7)$$

where σ is the maximum geodesic distance from a vertex within R_x to $\Phi_{L,\ell}^K(x)$.

Therefore, for a remaining LOI, we define its corresponding target vertex as the vertex with the highest weighted sum of the geometric and texture scores:

$$\Phi(x) = \arg \max_{v' \in R_x} \{w_1 \cdot C_L(x, v') + w_2 \cdot C_{L,\varepsilon}(x, v')\}, \quad (8)$$

where w_1 and w_2 are the weights for combining the scores.

3 Evaluation and Discussion

3.1 Dataset

We evaluated our methods on two datasets. The first dataset is from Skin3D [26] (annotated 3DBodyTex [18,19]). The second dataset comes from a 2D Imaging-Rich Total Body Photography system (IRTBP), from which the 3D textured meshes are derived from photogrammetry 3D reconstruction. The number of vertices is on average 300K and 600K for Skin3D and IRTBP datasets respectively. The runtime using 10 iterations is several minutes (on average) on an Intel i7-11857G7 processor. Example data of the two datasets can be found in the supplement.

3.2 Correspondence Localization Error and Success Rate

Average correspondence localization error (CLE) for individual subjects, defined as the geodesic distance between the ground-truth and the estimated lesion correspondence, is shown in Fig. 3. To interpret CLE in a clinical application, the localized correspondence is successful if its CLE is less than a threshold criterion. We measured the success rate as the percentage of the correctly localized skin lesions over the total number of skin lesion pairs in the dataset.

To compare our result to the existing method [26], we compute our success rates with the threshold criterion at 10 mm. As shown in table 1, the performance of our method is comparable to the previously reported longitudinal accuracy. The qualitative result of the localized correspondence in the Skin3D dataset is shown in Fig. 4. A table of parameters used in the experiments can be found in the supplement.




Fig. 3. Mean and standard deviation of correspondence localization error (CLE) for individual subjects in (a) Skin3D and (b) IRTBP datasets evaluated with the proposed methods. There are 10 and 3 subjects in the Skin3D and IRTBP datasets respectively. The number of LOIs for individual subjects is included in the parentheses.

Table 1. Comparison of the success rate on Skin3D dataset. Each metric is computed on a pair of meshes (for one subject) and averaged across paired meshes with the standard deviation shown in brackets. The method *Texture radius 50* and *Combined radius 50* are defined in Fig. 5.

	Skin3D [26]	Texture radius 50	Combined radius 50	Iterative algorithm
Success rate	0.5 (0.38)	0.48 (0.17)	0.45 (0.14)	0.57 (0.14)

3.3 Usage of Texture on 3D Surface

We believe that the geometric descriptor only provides a coarse correspondence while the local texture is more discriminating. Fig. 5 shows the success rate under different threshold criteria for the proposed methods. Since we have two combinations of defining source and target for two scans, we measured the result in both to ensure consistency (Fig. 5 (a) and (b)). As expected, we observed that using geometric information with body landmarks and geodesic distances is insufficient for localizing lesion correspondence accurately. However, the correspondence map $\Phi_{L,\epsilon}$ refined with local texture may lead to correspondences with large CLE when using a large region R_x . The figure shows the discriminating power and the large-error-prone property of using $\Phi_{L,\epsilon}$ to localize lesion correspondence with one iteration (relatively high success rates under strict criteria and relatively low success rates under loose criteria, compared to the correspondence map combining geometric and texture scores in Eqn. 8). The figure also shows the effectiveness of the proposed algorithm when the iterative anchor mechanism is used to localize lesion correspondence, having consistently higher success rates with the criteria within 20 mm.




Fig. 4. The qualitative result on Skin3D dataset. (a) shows a source scan (left) and the scan with annotated LOIs (right). (b) shows a target scan (left) and the scan with annotated LOIs (transparent sphere) and the estimated correspondence of LOIs (solid dot) (right). The lesion correspondence pairs are shown in the same color. Large CLE occurs when the local texture is not well-preserved across scans.

4 Conclusions and Limitations

The evolution of a skin lesion is an important sign of a potentially cancerous growth and total body photography is useful to keep track of skin lesions longitudinally. We proposed a novel framework that leverages geometric and texture information to effectively find lesion correspondence across TBP scans. The framework is evaluated on a private dataset and a public dataset with success rates that are comparable to those of the state-of-the-art method.

The proposed method assumes that the local texture enclosing the lesion and its surroundings should be similar from scan to scan. This may not hold when the appearance of the lesion changes dramatically (e.g. if the person acquires a tattoo). Also, the resolution of the mesh affects the precision of the positions of landmarks and lesions. In addition, the method may not work well with longitudinal data that has non-isometric deformation due to huge variations in body shape, inconsistent 3D reconstruction, or a dramatic change in pose and, therefore, topology, such as an open armpit versus a closed one.

In the future, the method needs to be evaluated on longitudinal data with longer duration and new lesions absent in the target. In addition, an automatic method to determine accurate landmarks is desirable. Note that although we rely on the manual selection of landmarks, the framework is still preferable over manually annotating lesion correspondences when a subject has hundreds of lesions. As the 3D capture of the full body becomes more prevalent with better quality in TBP, we expect that the proposed method will serve as a valuable step for the longitudinal tracking of skin lesions.




Fig. 5. Success rate under different criteria for the proposed algorithm on the Skin3D dataset. (a) and (b) show the results of two combinations of source and target meshes. *Iterative algorithm* is the proposed algorithm with the anchor mechanism. *Shape* is the method using $\Phi_{L,\ell}^1$. *Texture radius 25* and *texture radius 50* are the methods using $\Phi_{L,\epsilon}^1$ with ϵ_1 (Eqn. 3) selected at 25 mm and 50 mm. *Combined radius 25* and *combined radius 50* are the methods using Eqn. 8 with one iteration.

Acknowledgments The research was in part supported by the Intramural Research Program (IRP) of the NIH/NICHHD, Phase I of NSF STTR grant 2127051, and Phase I NIH/NIBIB STTR grant R41EB032304.

References

1. Abbasi, N.R., Shaw, H.M., Rigel, D.S., Friedman, R.J., McCarthy, W.H., Osman, I., Kopf, A.W., Polsky, D.: Early diagnosis of cutaneous melanoma: revisiting the abcd criteria. *Jama* **292**(22), 2771–2776 (2004)
2. Bogo, F., Romero, J., Peserico, E., Black, M.J.: Automated detection of new or evolving melanocytic lesions using a 3d body model. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 593–600. Springer (2014)
3. Datar, M., Lyu, I., Kim, S., Cates, J., Styner, M.A., Whitaker, R.: Geodesic distances to landmarks for dense correspondence on ensembles of complex shapes. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013: 16th International Conference, Nagoya, Japan, September 22–26, 2013, Proceedings, Part II 16. pp. 19–26. Springer (2013)
4. Guo, Y., Sohel, F., Bennamoun, M., Lu, M., Wan, J.: Rotational projection statistics for 3d local surface description and object recognition. *International journal of computer vision* **105**(1), 63–86 (2013)
5. Halpern, A.C.: Total body skin imaging as an aid to melanoma detection. In: Seminars in cutaneous medicine and surgery. vol. 22, pp. 2–8 (2003)
6. Kim, H., Kim, J., Kam, J., Park, J., Lee, S.: Deep virtual markers for articulated 3d shapes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11615–11625 (2021)
7. Korotkov, K., Quintana, J., Campos, R., Jesús-Silva, A., Iglesias, P., Puig, S., Malvehy, J., Garcia, R.: An improved skin lesion matching scheme in total body

- photography. *IEEE journal of biomedical and health informatics* **23**(2), 586–598 (2018)
8. Korotkov, K., Quintana, J., Puig, S., Malvey, J., Garcia, R.: A new total body scanning system for automatic change detection in multiple pigmented skin lesions. *IEEE transactions on medical imaging* **34**(1), 317–338 (2014)
 9. Li, Y., Esteva, A., Kuprel, B., Novoa, R., Ko, J., Thrun, S.: Skin cancer detection and tracking using data synthesis and deep learning. *arXiv preprint arXiv:1612.01074* (2016)
 10. Mcgregor, B.: Automatic registration of images of pigmented skin lesions. *Pattern Recognition* **31**(6), 805–817 (1998)
 11. Mirzaalian, H., Hamarneh, G., Lee, T.K.: A graph-based approach to skin mole matching incorporating template-normalized coordinates. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2152–2159. IEEE (2009)
 12. Mirzaalian, H., Lee, T.K., Hamarneh, G.: Uncertainty-based feature learning for skin lesion matching using a high order mrf optimization framework. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 98–105. Springer (2012)
 13. Mirzaalian, H., Lee, T.K., Hamarneh, G.: Spatial normalization of human back images for dermatological studies. *IEEE journal of biomedical and health informatics* **18**(4), 1494–1501 (2013)
 14. Mirzaalian, H., Lee, T.K., Hamarneh, G.: Skin lesion tracking using structured graphical models. *Medical image analysis* **27**, 84–92 (2016)
 15. Mitchel, T.W., Rusinkiewicz, S., Chirikjian, G.S., Kazhdan, M.: Echo: Extended convolution histogram of orientations for local surface description. In: *Computer Graphics Forum*. vol. 40, pp. 180–194. Wiley Online Library (2021)
 16. Perednia, D.A., White, R.G.: Automatic registration of multiple skin lesions by use of point pattern matching. *Computerized medical imaging and graphics* **16**(3), 205–216 (1992)
 17. Roning, J., Riech, M.: Registration of nevi in successive skin images for early detection of melanoma. In: *Proceedings. Fourteenth International Conference on Pattern Recognition (Cat. No. 98EX170)*. vol. 1, pp. 352–357. IEEE (1998)
 18. Saint, A., Ahmed, E., Cherenkova, K., Gusev, G., Aouada, D., Ottersten, B., et al.: 3dbodytex: Textured 3d body dataset. In: *2018 International Conference on 3D Vision (3DV)*. pp. 495–504. IEEE (2018)
 19. Saint, A., Cherenkova, K., Gusev, G., Aouada, D., Ottersten, B., et al.: Bodyfitr: robust automatic 3d human body fitting. In: *2019 IEEE International Conference on Image Processing (ICIP)*. pp. 484–488. IEEE (2019)
 20. Salti, S., Tombari, F., Di Stefano, L.: Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding* **125**, 251–264 (2014)
 21. Strakowska, M., Kociołek, M.: Skin lesion matching algorithm for application in full body imaging systems. In: *International Conference on Information Technologies in Biomedicine*. pp. 222–233. Springer (2022)
 22. Strzelecki, M.H., Strakowska, M., Kozłowski, M., Urbańczyk, T., Wielowieyska-Szybińska, D., Kociołek, M.: Skin lesion detection algorithms in whole body images. *Sensors* **21**(19), 6639 (2021)
 23. Tombari, F., Salti, S., Di Stefano, L.: Unique signatures of histograms for local surface description. In: *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part III 11*. pp. 356–369. Springer (2010)

24. Tombari, F., Salti, S., Di Stefano, L.: A combined texture-shape descriptor for enhanced 3d feature matching. In: 2011 18th IEEE international conference on image processing. pp. 809–812. IEEE (2011)
25. White, R.G., Perednia, D.A.: Automatic derivation of initial match points for paired digital images of skin. *Computerized medical imaging and graphics* **16**(3), 217–225 (1992)
26. Zhao, M., Kawahara, J., Abhishek, K., Shamanian, S., Hamarneh, G.: Skin3d: Detection and longitudinal tracking of pigmented skin lesions in 3d total-body textured meshes. *Medical Image Analysis* **77**, 102329 (2022)