



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

S-SYNTH: Knowledge-Based, Synthetic Generation of Skin Images

Andrea Kim*, Niloufar Saharkhiz*, Elena Sizikova*, Miguel Lago, Berkman Sahiner, Jana Delfino, and Aldo Badano

Office of Science and Engineering Laboratories
Center for Devices and Radiological Health
U.S. Food and Drug Administration
Silver Spring, MD 20993 USA
elena.sizikova@fda.hhs.gov

*-These authors contributed equally to this work.

Abstract. Development of artificial intelligence (AI) techniques in medical imaging requires access to large-scale and diverse datasets for training and evaluation. In dermatology, obtaining such datasets remains challenging due to significant variations in patient populations, illumination conditions, and acquisition system characteristics. In this work, we propose S-SYNTH, the first knowledge-based, adaptable open-source skin simulation framework to rapidly generate synthetic skin, 3D models and digitally rendered images, using an anatomically inspired multi-layer, multi-component skin and growing lesion model. The skin model allows for controlled variation in skin appearance, such as skin color, presence of hair, lesion shape, and blood fraction among other parameters. We use this framework to study the effect of possible variations on the development and evaluation of AI models for skin lesion segmentation, and show that results obtained using synthetic data follow similar comparative trends as real dermatologic images, while mitigating biases and limitations from existing datasets including small dataset size, lack of diversity, and under-representation.

Keywords: Synthetic Data · Simulation · Dermatology

1 Introduction

Robust and generalizable artificial intelligence (AI) for medical applications requires large datasets representative of intended populations and their subgroups, relying on a time-consuming and laborious annotation process by clinical experts (e.g., labeling pixel-level segmentation masks or identifying findings). In this work, we explore the potential of a skin simulation model and the practical utility of resulting synthetic images in development and evaluation of AI-based techniques.

Segmenting and labeling dermatologic images is time-consuming and challenging, and thus public dermatologic datasets generally contain few examples

and may not equitably represent intended patient populations [11,17]. For instance, the majority of the publicly available dermatologic images are associated with lighter skin populations, and dark skin tones are typically under-represented [23]. This class imbalance has shown to negatively affect darker skin tones in skin lesion segmentation tasks [8]. Such challenges may introduce biases in training and evaluation of medical AI devices for dermatologic applications. Moreover, real skin datasets are prone to inaccurate annotations [3], such as labeling errors, near duplicates, and incorrect labels, as well as confounders [36], such as rulers, dark borders, dense hairs and air pockets, which may negatively affect AI-based diagnostic tools and any data-driven methods ([29,30]) used to generate synthetic data.

To address these limitations, we propose to simulate dermatologic images with a knowledge-based approach [5], where properties of skin and lesions, determined via detailed physics models, along with rendering conditions are systematically varied. We created S-SYNTH, an open-source skin simulation framework for generating synthetic dermatologic images, and show that subgroup analysis using S-SYNTH images closely mimics comparative performance trends in real-patient skin datasets, while mitigating effects of mislabeled examples and limitations of a limited patient evaluation dataset. Although the known confounders [36] were taken into account to generate the images, our model did not propagate the confounding factors to the synthetic data, unlike data-driven models ([29,30]). To our knowledge, S-SYNTH is the first to use knowledge-based approach to produce skin images with and without lesions. Our contributions can be summarized as follows:

- We describe S-SYNTH, an open-source, flexible framework for creation of highly-detailed 3D skin models and digitally rendered synthetic images of diverse human skin tones, with full control of underlying parameters and the image formation process¹.
- We systematically evaluate S-SYNTH synthetic images for training and testing applications. Specifically, we show that S-SYNTH synthetic images improve segmentation performance when only a limited set of real images is available for training. We also show comparative trends between S-SYNTH synthetic images and real-patient examples (according to skin color and lesion size) are similar.

2 Related Work

A particular challenge in practical development and evaluation of dermatologic AI is the lack of labeled datasets. Of the few public labeled datasets that exist, only a fraction have fine grained annotations, such as pixel-level segmentation levels for lesions [25] or hair [19]. The annotation process is error-prone and can have large variations across truthers (particularly when datasets are collected

¹ Code and supporting data are available at: <https://github.com/DIDSR/ssynth-release>

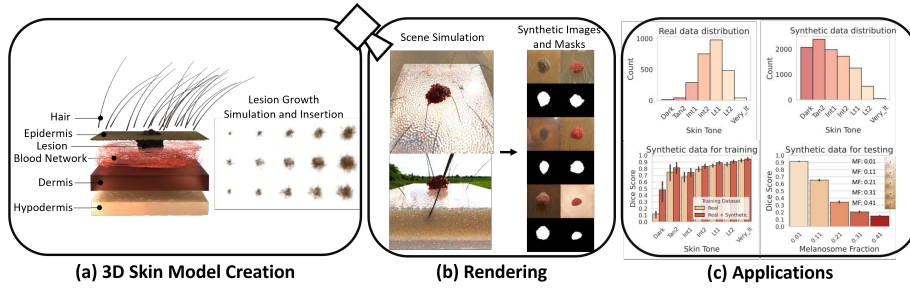


Fig. 1. Overview of (a) a digital skin model generated in Houdini, and an example projection of 3 different skin lesion volumes at 5 growing time points, (b) digital rendering, with examples of generated synthetic images and their corresponding lesion masks, and (c) distribution of skin tone for real-patient images (ISIC) and synthetic images, as well as application of synthetic images for training and testing of an AI device for a skin lesion segmentation task.

for different purposes and no consistent labeling protocol is established). In part due to these annotation difficulties, there is a lack of transparency and potential bias within the datasets used for dermatologic AI studies. A recent investigation found that from a total of 70 unique studies that involved AI models for dermoscopic imaging, only about 7 (10%) included skin tone information for at least one dataset used [14]. Synthetically generated data has been explored to address these challenges. Generative models have been used to create synthetic skin lesions to tackle the class imbalance problem and improve the performance of the skin lesion classifiers [6,7,16] or enhance lesion segmentation [12,26]. Rezk et al. have shown that synthetic skin images generated using style transfer and deep blending improve skin lesion classification accuracy by increasing skin color diversity [27]. More recently, diffusion models have been used to generate synthetic images to improve skin disease classifier performance [29,30]. Nevertheless, use of AI to generate synthetic images is limited by the requirement of constant supervision and adjustment of hyperparameters [6,7], presence of artifacts, such as skin hair or lighting effects [26], or low prevalence of examples with dark skin tones within the original training dataset, and a lack of ability to represent all physical and morphological features of disease [27] and image acquisition systems (see [25] for more information).

3 S-SYNTH: Skin Simulation Framework

Our approach for generating synthetic skin involves the construction of a 3D digital object model comprising of skin tissue (epidermis, dermis, hypodermis), blood network, hair, and a lesion, as shown in Fig. 1a. This process is implemented and automated in Houdini [2], a software for 3D modeling, animation, and visual effects, via a Python API. Once created, each model is subsequently

processed through a research-oriented rendering system, Mitsuba 3 [21], to generate each synthetic rendering (Fig 1b). The approach can be automated to create large databases with controlled variation for various AI analysis applications (Fig 1c). In the following sections, we describe the digital skin model and rendering process.

3.1 Digital Skin Model

Skin Tissue Each skin sample is created from a multi-layer model, with each layer representing a component of skin tissue. Thicknesses of the topmost two layers are based on values reported in [35]: epidermis (20-150 μm) and dermis (1-4 mm). To introduce geometric variability, we incorporated surface roughness (noise) into the top surfaces of epidermis, hypodermis, and papillary dermis (located between the epidermis and dermis). Roughness values were randomized within predefined constraints, enhancing the naturalistic appearance and diversity of the generated skin models.

Blood Network and Hair The blood network model is created by solving a shortest path problem on randomly distributed points within a tetrahedral mesh generated from a primitive cube. Points from the bottom of the cube, corresponding to the lower blood network capillaries, were designated as start points, while those from top of the cube, representing the upper capillaries, were designated as end points. To generate each blood network model, the numbers and positions of starting and ending points were randomized. Hair models were constructed by manipulating hair properties (density, length, distribution, thickness, and curvature) in Houdini to generate a diverse range of hair appearances.

Growing Lesion Model We developed a probabilistic, growing lesion model that generates a 3D volume with stochastic growth of lesion shape and size, based on [31]. Growth starts with one active cell (i.e., voxel). We define *active cells* as those cells that exist and can grow at a given time point. Within each time point, we iterated over the active cells c and selected them for growth based on directional probabilities (see Fig. 2). To control the growing direction, the probability of an active cell to grow outside the skin is set to extremely low, while the probability for inwards growing is set to be relatively high. Probabilities were based on a Gaussian distribution G , and are updated in each time point for added variability.

To control for lesion shape irregularity, we set an irregularity probability for each cell C_p , which starts C_i number of recursive iterations of the same growing algorithm on that cell’s location, independent of the regular growth. These cells can themselves trigger another recursive growth up to a maximum of C_r recursions. We controlled the generation parameters to create both regular ($C_p = 0.0001$) and irregular lesions ($C_p = 0.001$). Sample images of lesions and growth evolution can be seen in Fig 1. Additional details can be found in the Supplementary Material. Once created, the 3D lesion volume is inserted into the skin tissue model (see [22] for more details).




Fig. 2. Growing direction probabilities for each cell c in the 3D plane. Note that values are zeros in the *outwards* direction, to avoid the lesion growth outside the skin.

3.2 Rendering and Image Formation

Optical Material Properties Melanosomes are the primary contributors to optical absorption in the epidermis. In the dermis, optical absorption is mainly attributed to blood, while, in the hypodermis, lipid and fat components are the predominant sources of optical absorption. Thus, each skin layer is assigned an optical material, containing the index of refraction (IOR) and the spectral distribution of the absorption and scattering coefficients. Following the methodology established by Jacques et al. [20], we calculated the optical absorption properties for each the four distinct tissue types: epidermis, dermis, hypodermis, and blood. We model reduced scattering of tissue using $\mu'_s = a(\lambda/500)^{-b}$, where λ is the light wavelength [9]. Finally, for each skin layer material, we define a surface scattering model using a bidirectional scattering distribution function (BSDF), based on [34].

Lighting and Camera To account for lighting variations, we render each skin model using a collection of High Dynamic Range Imaging (HDRI) images capturing various environmental lighting conditions from [1]. The Volumetric Path Tracer is a rendering algorithm that simulates the paths taken by light as it interacts with the 3D scene. It is particularly effective for scenes with volumetric effects or materials with complex light interactions. Spectral multiple importance sampling is a technique used to address the spectral nature of light. In scenes with materials that exhibit variations in spectral properties (e.g., spectral absorption), this method accounts for the complex interactions of light with materials across different wavelengths, such as those encountered in biological tissues. A perspective sensor with Mitsuba’s HDR Film is placed 1.5cm above the skin with a field of view (FOV) of 75 degrees, facing perpendicularly to the top surface of the epidermis. The Integrator (Mitsuba 3’s plugin for the rendering technique) is set to Volumetric Path Tracer with spectral multiple importance sampling. We chose this specific configuration due to its ability to simulate spectrally-varying optical properties across varying skin tones in simulated lighting conditions. We use 124 samples per pixel (SPP) and render images at resolution 1024x1024 pixels. Each image and mask take about three minutes to generate on a GPU. Sample generated images can be seen in Fig. 3.

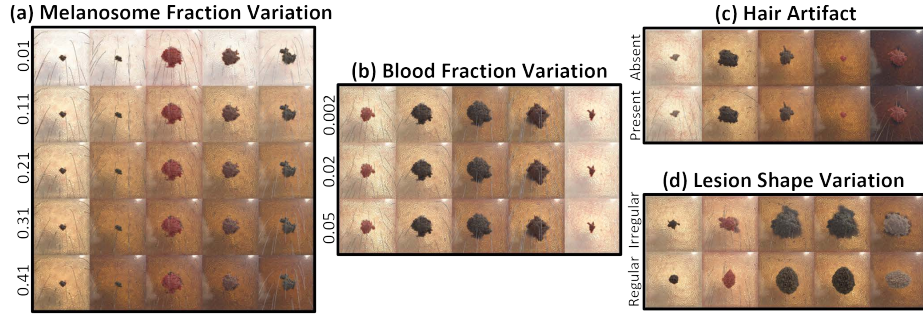


Fig. 3. Examples of S-SYNTH images generated with variations of (a) melanosome fraction, (b) blood fraction, (c) hair artifact, and (d) lesion shape.

3.3 AI Device Description

To demonstrate the usage of S-SYNTH for augmenting real patient data, we experiment with real and synthetic examples on the task of skin lesion segmentation, which is important for timely treatment decisions [28]. We rely on DermoSegDiff [10], a state of the art diffusion-based skin segmentation model, to differentiate lesions from the background region. We used the “dsd_i01” configuration, with $\text{dim}_x=16$, $\text{dim}_g=8$, and default training parameters. Images were segmented at resolution 128×128 , and results were evaluated with Dice coefficient. We trained one model per experiment and reported standard deviation over test images.

Datasets For real-patient datasets, we used two publicly available skin lesion segmentation datasets: ISIC18 [13] (ISIC) and HAM10K [32] (HAM). The same pre-processing and data splits described in [13,4] were applied, where the ISIC and HAM images were divided into (7200, 1800, 1015) and (1815, 259, 520) for the baseline model (training, validation, test), respectively. We used the individual typology angle (ITA) metric [23] to estimate skin tone from the non-lesion area. Finally, to better match lesion sizes across real and synthetic examples, S-SYNTH images were cropped to ensure a variety of lesion sizes. With S-SYNTH, we generated 10,000 images with randomized model properties and lighting, and subsequently 19,965 testing images with controlled variation: 5,445 with blood variation, 9,075 with melanosome fraction variation, 3,630 with different lesion regularity, and 1,815 without hair.

4 Results and Discussion

We present segmentation performance based on three aspects: impact of training set composition comprising different ratios of real and synthetic images, impact of various physiological or rendering parameters of the synthetic images when

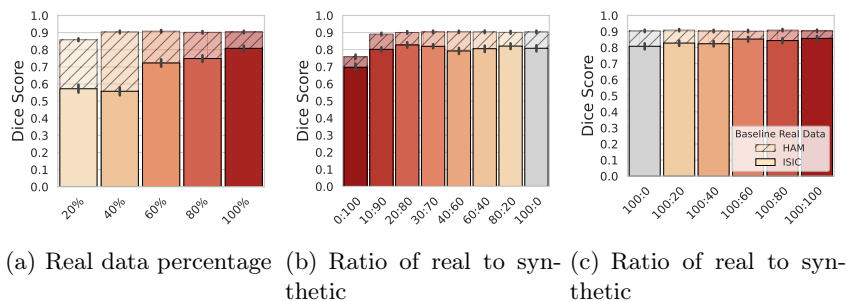


Fig. 4. Model performance changes when the training data is composed of (a) different numbers of the real images, (b) different proportions of real images replaced with synthetic images (c) different proportions of synthetic images added to real images.

used as test set, and impact of specific characteristics of synthetic images that are measurable on real images.

Synthetic Data in Training We systematically evaluated the effect of the training data composition on the real test set performance (similar to the popular Train-Synthetic-Test-Real (TSTR) protocol [15]) on both ISIC and HAM. As shown in Fig. 4a, performance improved when more real examples are available for training. When real images are replaced with synthetic images, the resulting performance is comparable to that of the baseline model trained on full ISIC (or HAM) training set (Fig. 4b). However, performance dropped when the training set was fully composed of synthetic images, presumably due to the domain shift between the real and synthetic images [18]. More importantly, each model that was trained on a particular subset of the ISIC (or HAM) and supplemented with synthetic images resulted in a higher Dice score than the model that was trained on the same subset of only the real images. Finally, when full training sets were supplemented with synthetic images (Fig. 4c), performance improved for ISIC, but saturated for HAM, likely because ISIC was too small to learn all modes of variation, which synthetic data helped address. We also found that performance for darker skin tones improved with the addition of synthetic examples (see Supplementary Material).

Synthetic Data for Testing To check whether synthetic images can produce meaningful performance trends, we evaluated models trained on real examples using synthetic test images. Performance dropped for increasing blood fraction (Fig. 5a), and increasing melanosome fraction (Fig. 5b) of the S-SYNTH images. Higher melanosome fraction corresponds to darker skin, and the corresponding performance drop is a known bias in dermatologic AI [27]. Models tested on lesions with regular shape exhibited higher performance compared to the ones

tested on irregular lesions (Fig. 5c). Similar to [24], we found that presence of hair slightly decreased performance (Fig. 5d).

To assess the similarity between the real and synthetic images in measurable characteristics, we evaluated performance changes as a function of skin color (ITA), lesion circularity, and lesion relative area. Although performance on synthetic images was lower compared to real images, similar trends were observed. Specifically, performance in both synthetic and real examples increased with lighter skin (Fig. 6a). As illustrated in (Fig. 6b), the impact of lesion circularity on model performance showed similar trends within the circularity range shared between the real and synthetic images, where more circular images were easier to segment. Finally, performance dropped with larger relative lesion areas (Fig. 6c), possibly due to less precise delineations in larger lesions in the training patient data.




Fig. 5. Model performance when trained on real-patient datasets (ISIC, HAM) and tested on synthetic (S-SYNTH) images generated with different parameters. BF: blood fraction, MF: melanosome fraction.




Fig. 6. Performance as a function of: (a) skin tone, (b) lesion circularity, and (c) relative lesion area for models trained on HAM dataset and tested on real (HAM, ISIC) or synthetic (S-SYNTH) images.

Limitations S-SYNTH is the first skin simulation framework that generates synthetic skin images with controllable variations. There are a number of limitations to our work. First, S-SYNTH does not model any specific disease presentation, and additional work needs to be performed to evaluate the realism of the generated examples. Second, the current rendering techniques utilizes an RGB camera setup, and additional modifications allowing for multispectral imaging, when validated, might be useful for more advanced algorithms that use a broader variety of wavelengths. Lastly, although our lesion growth model surpasses those based on simple geometrical shapes [33], it is to some degree simplistic, particularly compared to real anatomical and pathological skin characteristics. The current lesion growth model varies along the simulation timeline and is affected by material changes but can be further improved by considering the interaction between the lesion and non-lesion areas.

5 Conclusion and Future Work

We presented S-SYNTH, a novel, knowledge-based simulation pipeline for synthetic generation of dermoscopic images and skin lesions. S-SYNTH procedurally generates multi-layer 3D skin models, with consideration to optical skin characteristics, and digitally renders synthetic images under realistic lighting conditions. Using S-SYNTH, we generated realistic and varied skin surface models for sub-surface scattering simulation in the context of skin lesion segmentation, and compared the resulting synthetic images to public dermatologic benchmarks. We showed that examples created using S-SYNTH can be used to augment limited real datasets and identify performance trends.

Acknowledgments. We thank anonymous reviewers for helpful suggestions. We thank the OpenHPC and RST teams (OSEL/CDRH/FDA) for providing help with experiments and data release. This is a contribution of the US Food and Drug Administration and is not subject to copyright. The mention of commercial products herein is not to be construed as either an actual or implied endorsement of such products by the Department of Health and Human Services.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Poly Haven: the public 3D asset library. <https://polyhaven.com/hdris>
2. Sidefx: Houdini (2022), <https://www.sidefx.com/docs/houdini/index.html>
3. Abhishek, K., Jain, A., Hamarneh, G.: Investigating the quality of dermamnist and fitzpatrick17k dermatological image datasets. arXiv preprint arXiv:2401.14497 (2024)
4. Azad, R., Al-Antary, M.T., Heidari, M., Merhof, D.: Transnorm: Transformer provides a strong spatial normalization mechanism for a deep segmentation model. IEEE Access (2022)

5. Badano, A., Lago, M., Sizikova, E., Delfino, J.G., Guan, S., Anastasio, M.A., Sahiner, B.: The stochastic digital human is now enrolling for in silico imaging trials—methods and tools for generating digital cohorts. *Progress in Biomedical Engineering* (2023)
6. Baur, C., Albarqouni, S., Navab, N.: Melanogans: high resolution skin lesion synthesis with gans. arXiv:1804.04338 (2018)
7. Behara, K., Bhero, E., Agee, J.T.: Skin lesion synthesis and classification using an improved dcgan classifier. *Diagnostics* (2023)
8. Benčević, M., Habijan, M., Galić, I., Babin, D., Pižurica, A.: Understanding skin color bias in deep learning-based skin lesion segmentation. *Computer Methods and Programs in Biomedicine* (2024)
9. Bench, C., Hauptmann, A., Cox, B.: Toward accurate quantitative photoacoustic imaging: learning vascular blood oxygen saturation in three dimensions. *Journal of Biomedical Optics* **25**(8), 085003–085003 (2020)
10. Bozorgpour, A., Sadegheih, Y., Kazerouni, A., Azad, R., Merhof, D.: Dermosegdiff: A boundary-aware segmentation diffusion model for skin lesion delineation. In: *International Workshop on PRedictive Intelligence In MEDicine (PRIME)*. Springer (2023)
11. Chen, M.L., Rotemberg, V., Lester, J.C., Novoa, R.A., Chiou, A.S., Daneshjou, R.: Evaluation of diagnosis diversity in artificial intelligence datasets: a scoping review. *British Journal of Dermatology* (2023)
12. Chi, Y., Bi, L., Kim, J., Feng, D., Kumar, A.: Controlled synthesis of dermoscopic images via a new color labeled generative style transfer network to enhance melanoma segmentation. In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE (2018)
13. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M.E., Dusza, S., Gutman, D., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M., et al.: Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC). arXiv:1902.03368 (2019)
14. Daneshjou, R., Smith, M.P., Sun, M.D., Rotemberg, V., Zou, J.: Lack of transparency and potential bias in artificial intelligence data sets and algorithms: a scoping review. *JAMA Dermatology* (2021)
15. Esteban, C., Hyland, S.L., Rätsch, G.: Real-valued (medical) time series generation with recurrent conditional gans. arXiv:1706.02633 (2017)
16. Ghorbani, A., Natarajan, V., Coz, D., Liu, Y.: Dermgan: Synthetic generation of clinical skin images with pathology. In: *Machine Learning for Health (ML4H)*. PMLR (2020)
17. Gröger, F., Lionetti, S., Gottfrois, P., Gonzalez-Jimenez, A., Groh, M., Daneshjou, R., Navarini, A.A., Pouly, M., Consortium, L., et al.: Towards reliable dermatology evaluation benchmarks. In: *Machine Learning for Health (ML4H)*. PMLR (2023)
18. Guan, H., Liu, M.: Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering* (2021)
19. Hossain, S.I., Roy, S.S., De Herve, J.D.G., Mercer, R.E., Nguifo, E.M.: A skin lesion hair mask dataset with fine-grained annotations. *Data in Brief* (2023)
20. Jacques, S.L.: Optical properties of biological tissues: a review. *Physics in Medicine & Biology* (2013)
21. Jakob, W., Speierer, S., Roussel, N., Nimier-David, M., Vicini, D., Zeltner, T., Nicolet, B., Crespo, M., Leroy, V., Zhang, Z.: Mitsuba 3 renderer (2022), <https://mitsuba-renderer.org>

22. Kim, A.S., Sengupta, A., Badano, A.: Automated animation pipeline for visualizing in silico tumor growth models. In: *Medical Imaging 2023: Physics of Medical Imaging*. SPIE (2023)
23. Kinyanjui, N.M., Odonga, T., Cintas, C., Codella, N.C., Panda, R., Sattigeri, P., Varshney, K.R.: Estimating skin tone and effects on classification performance in dermatology datasets. arXiv:1910.13268 (2019)
24. Li, W., Raj, A.N.J., Tjahjadi, T., Zhuang, Z.: Digital hair removal by deep learning for skin lesion segmentation. *Pattern Recognition* (2021)
25. Mirikharaji, Z., Abhishek, K., Bissoto, A., Barata, C., Avila, S., Valle, E., Celebi, M.E., Hamarneh, G.: A survey on deep learning for skin lesion segmentation. *Medical Image Analysis* (2023)
26. Oliveira, D.A.B.: Controllable skin lesion synthesis using texture patches, bézier curves and conditional gans. In: *International Symposium on Biomedical Imaging (ISBI)*. IEEE (2020)
27. Rezk, E., Eltorki, M., El-Dakhakhni, W., et al.: Improving skin color diversity in cancer detection: deep learning approach. *JMIR Dermatology* (2022)
28. Ruan, J., Xie, M., Gao, J., Liu, T., Fu, Y.: Ege-unet: an efficient group enhanced unet for skin lesion segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer (2023)
29. Sagers, L.W., Diao, J.A., Groh, M., Rajpurkar, P., Adamson, A.S., Manrai, A.K.: Improving dermatology classifiers across populations using images generated by large diffusion models. arXiv:2211.13352 (2022)
30. Sagers, L.W., Diao, J.A., Melas-Kyriazi, L., Groh, M., Rajpurkar, P., Adamson, A.S., Rotemberg, V., Daneshjou, R., Manrai, A.K.: Augmenting medical image classifiers with synthetic data from latent diffusion models. arXiv:2308.12453 (2023)
31. Sengupta, A., Sharma, D., Badano, A.: Computational model of tumor growth for in silico trials. In: *Medical Imaging 2021: Physics of Medical Imaging*. SPIE (2021)
32. Tschandl, P., Rosendahl, C., Kittler, H.: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *scientific data*. 2018; 5: 180161. Search in (2018)
33. Vasudev, V., Piepers, B., Maidment, A.D., Kimpe, T., Platisa, L., Philips, W., Bakic, P.R.: Simulation pipeline for virtual clinical trials of dermatology images. In: *Medical Imaging 2019: Physics of Medical Imaging*. vol. 10948, pp. 600–609. SPIE (2019)
34. Walter, B., Marschner, S.R., Li, H., Torrance, K.E.: Microfacet models for refraction through rough surfaces. In: *Eurographics Conference on Rendering Techniques* (2007)
35. Walters, K.A., Roberts, M.S.: The structure and function of skin. In: *Dermatological and Transdermal Formulations*. CRC press (2002)
36. Yan, S., Yu, Z., Zhang, X., Mahapatra, D., Chandra, S.S., Janda, M., Soyer, P., Ge, Z.: Towards trustable skin cancer diagnosis via rewriting model’s decision. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11568–11577 (2023)