

Physics-Guided Synthetic High-Frequency Ultrasound Generation for Skin Layer Segmentation

Junkyung Ju¹, Kyungho Yoon², and Minwoo Shin^{1*}

¹ Department of Software, Yonsei University (Mirae Campus), Wonju, 26493, Republic of Korea

² School of Mathematics and Computing, Yonsei University, Seoul, 03722, Republic of Korea
`mshin@yonsei.ac.kr`

Abstract. High-frequency ultrasound (HFUS) enables noninvasive visualization of superficial skin structures, but automated skin-layer analysis is limited by the scarcity of densely annotated data. Existing real HFUS datasets commonly provide annotations for superficial targets such as the epidermis and subepidermal low-echogenic band (SLEB), while dense labels for deeper structures such as dermis, subcutaneous tissue, fascia, and muscle are rarely available. We propose a physics-guided synthetic HFUS generation framework for skin layer segmentation. The framework constructs multilayer acoustic skin phantoms, assigns layer-dependent acoustic properties, and uses k-Wave simulation to generate paired synthetic HFUS images, dense layer masks, and simulation metadata. To evaluate whether the generated data provide transferable supervision, we use the generated synthetic data for downstream segmentation pretraining and fine-tune the models on a public real HFUS dataset. Synthetic pretraining followed by real fine-tuning achieved real-domain performance comparable to real-only training and improved mean Dice/IoU in three of the four trainable architectures evaluated. These results suggest that physics-guided synthetic HFUS images contain transferable anatomical and textural cues for real-domain skin layer segmentation, although further reduction of the synthetic-real appearance gap is needed to enable greater gains. The code and data are available at <https://github.com/Finn-02/synthetic-hfus-skin-layer-segmentation>.

Keywords: High-frequency ultrasound · Synthetic generation · Skin layer segmentation · SLEB · k-Wave

1 Introduction

High-frequency ultrasound (HFUS) is widely used in dermatologic imaging because it provides high-resolution visualization of superficial tissue structures [1,16].

* Corresponding author.

In skin analysis, localizing the epidermis, SLEB, dermis, subcutaneous tissue, fascia, and deeper tissue regions is important for measuring thickness, monitoring structural changes, and supporting future image-guided applications. However, automated HFUS skin analysis remains difficult because images contain speckle noise, weak layer contrast, attenuation, probe-dependent appearance variation, and ambiguous tissue boundaries [4,8,17].

A major bottleneck is the scarcity of densely annotated HFUS skin datasets. Existing public resources mainly provide annotations for the epidermis and SLEB [4,5,7], whereas dense labels for dermis, subcutaneous tissue, fascia, and muscle are rarely available. Thus, supervised HFUS segmentation studies are often restricted to a narrow subset of skin anatomy.

Synthetic data generation can alleviate this limitation because simulation-based data can provide exact masks and generation metadata. Prior work has highlighted simulation and synthesis for medical image analysis under limited annotation settings [10], and recent synthetic ultrasound studies show that generated images can support downstream learning [9,11,26]. Nevertheless, skin-specific HFUS generation remains underexplored: existing skin HFUS studies mainly focus on real-image classification, quality assessment, or segmentation rather than physics-guided generation of densely labeled multilayer skin images. A useful synthetic dataset should therefore preserve the anatomical ordering of skin layers while allowing controlled variation in thickness, interface morphology, and tissue texture. This motivates a generator that is driven by acoustic layer maps rather than by image-level style transfer alone.

We propose a physics-guided synthetic HFUS generation framework for skin layer segmentation. The framework constructs multilayer acoustic skin phantoms, generates synthetic HFUS images using k-Wave [27,28], and automatically exports dense 8-class masks and metadata. We then use downstream epidermis/SLEB segmentation to validate whether the generated data provide transferable supervision for real HFUS analysis.

Our contributions are as follows:

1. We propose a physics-guided synthetic HFUS image generation framework based on multilayer skin acoustic phantoms and k-Wave simulation.
2. We introduce controllable structural and appearance variations for layer-aware skin synthesis, producing dense labels for air, coupling medium, epidermis, SLEB, dermis, subcutaneous tissue, fascia, and muscle.
3. We evaluate downstream utility through real-domain epidermis/SLEB segmentation, showing that synthetic pretraining followed by real fine-tuning achieves performance comparable to real-only training and improves mean Dice/IoU in three of the four trainable architectures.

2 Related Work

HFUS skin analysis and segmentation. HFUS has been used to assess epidermal thickness, dermal echogenicity, SLEB-related changes, inflammatory dermatoses, skin aging, and aesthetic treatment effects [1,16,20,29]. Learning-based

HFUS studies have addressed image classification, quality assessment, epidermis segmentation, and epidermis/SLEB delineation [4,5,6,8], with recent work extending automated segmentation to entry echo, SLEB, and dermis [24]. These studies demonstrate the importance of automated HFUS skin analysis, but primarily operate on real images rather than physics-guided generation of densely labeled multilayer skin datasets.

Synthetic ultrasound generation and simulation. Data-driven ultrasound synthesis has been explored using diffusion and GAN-based models for echocardiography, breast ultrasound, and musculoskeletal ultrasound [9,11,26,31]. These methods show the value of synthetic ultrasound for downstream learning, but generally learn image distributions from data and do not explicitly construct multilayer skin acoustic maps with paired dense layer labels. Physics-based simulators provide a complementary route because the anatomical or acoustic model is known during image formation. Field II simulates ultrasound systems and transducer fields [12,13], while k-Wave performs time-domain acoustic simulation in heterogeneous media [27,28] and includes B-mode simulation examples [14]. Despite the availability of these simulation tools, recent simulator reviews, and learning-based acceleration methods for ultrasound acoustic simulation [22,23,25], physics-guided generation of densely labeled multilayer HFUS skin images for segmentation pretraining remains underexplored.

3 Method

3.1 Overview of Physics-Guided Synthetic HFUS Generation

Fig. 1 illustrates how the proposed simulator converts a multilayer acoustic skin phantom into paired synthetic HFUS images and dense segmentation masks. Starting from layer-dependent acoustic properties, the framework constructs density and sound-speed maps, performs k-Wave simulation, and exports both grayscale and green-channel synthetic HFUS images. Because the simulation is driven by known anatomical layer maps, the corresponding dense masks are obtained without manual annotation. This links acoustic phantom design, image generation, and pixel-level supervision within a single physics-guided pipeline. Each simulation output is stored together with its acoustic maps, raw class map, visual mask, overlay, and generation metadata. This organization makes the generation process traceable and allows downstream models to be trained on either the full multilayer label space or a task-specific label mapping.

3.2 Multilayer Acoustic Skin Phantom

The simulation pipeline uses k-Wave acoustic simulation to convert a multilayer acoustic skin phantom into synthetic HFUS images with raw class-ID masks [14,27,28]. The target label space contains eight classes: air, coupling medium, epidermis, SLEB, dermis, subcutaneous tissue, fascia, and muscle. Each class is assigned acoustic and structural properties through the simulator configuration. Table 1 summarizes the principal material parameters used to initialize




Fig. 1. Representative outputs of the proposed physics-guided synthetic HFUS generation framework. From left to right, the panels show the density map, sound-speed map, grayscale simulated ultrasound image, green-channel synthetic HFUS image, and dense layer mask. The acoustic property maps are simulated with k-Wave to generate ultrasound images, and the known anatomical layer map provides dense segmentation supervision.

the phantom. SLEB is implemented as a low-scatter inserted band inside the upper skin region; its listed sound speed and density are target values used during SLEB material mixing.

Table 1. Principal acoustic phantom parameters used in the synthetic HFUS generator. Sound-speed and density values were informed by reported acoustic properties of human skin, skin-mimicking phantoms, and soft tissues, while scatter levels are simulation controls for layer-dependent HFUS texture [2,15,18,19].

Layer	Sound speed c (m/s)	Density ρ (kg/m ³)	Base scatter std.
Air / near-field medium	1480	1000	2.0×10^{-5}
Coupling medium	1484	1002	2.0×10^{-4}
Epidermis	1628	1125	1.0×10^{-2}
SLEB	1518	1015	low-scatter band
Dermis	1578	1064	2.2×10^{-2}
Subcutaneous tissue	1556	1032	6.2×10^{-3}
Fascia	1512–1558	978–1036	thin-sheet controlled
Muscle	1560–1565	1040–1045	$4.5\text{--}7.2 \times 10^{-3}$

The simulator is designed to preserve thin superficial structures while extending the label space to deeper skin regions. The epidermis and SLEB are retained as separate classes. The dermis-to-subcutis transition is modeled as a gradual acoustic transition rather than a single artificial boundary, and fascia is represented as a thin physical layer at the subcutis–muscle interface. Muscle texture is modeled as a deeper anisotropic region rather than as a uniform background. This design allows the same synthetic image to support both the conventional epidermis/SLEB task and a broader multilayer segmentation task.

3.3 Structural and Appearance Variation

To avoid generating a single fixed skin geometry, the generator introduces three families of variation. First, layer-thickness variation changes the relative thickness of superficial and deeper layers across samples. Second, boundary-shape variation perturbs the morphology and smoothness of layer interfaces, allowing non-flat epidermal, SLEB, dermal, and subcutaneous boundaries. Third, tissue-heterogeneity variation modifies dermal scattering, subcutaneous hypoechoic texture, fascia visibility, and muscle texture. After generation, label-preserving image augmentation further increases appearance diversity while preserving the corresponding segmentation masks. We intentionally separate these controls: structural variation changes the simulated anatomy before image formation, whereas augmentation changes the exported image and mask pair after generation. This separation keeps anatomy-level diversity distinct from training-time appearance augmentation.

3.4 k-Wave Simulation, Export, and Task Mapping

The production workflow uses a 20 MHz simulation configuration with 96 scan lines and exports synthetic HFUS images at 2067×1555 pixels. This export size is intentionally matched to the publicly available real HFUS skin images used for comparison, which are also provided at 2067×1555 pixels [7]. The final field of view corresponds to a 7.0 mm axial depth. To match the green display format of the publicly available real HFUS images used for fine-tuning and evaluation, we use the fundamental green-channel rendering as the segmentation input rather than the grayscale rendering. The green-channel rendering is derived from the same simulated fundamental image and does not alter the paired segmentation labels. The raw integer mask independently preserves the class labels defined by the common class map.

The generated dataset is organized as paired images, masks, visual masks, overlays, and metadata. In the current dataset, 200 original synthetic samples are produced from parent simulations and paired variants. For real-domain epidermis/SLEB segmentation, the synthetic 8-class masks are mapped into a 3-class target space: other, epidermis, and SLEB. Specifically, the epidermis class is retained as epidermis, the SLEB class is retained as SLEB, and all remaining synthetic classes are mapped to the other class. This mapping allows the synthetic data to be used as anatomical pretraining supervision while the final evaluation remains directly comparable to the public Mendeley HFUS epidermis/SLEB annotations.

4 Experiments

4.1 Experiment 1: Generating the Synthetic Dataset

The first experiment evaluates whether the proposed simulator can produce diverse synthetic HFUS skin images with paired dense labels. We generate syn-

thetic samples by varying layer thickness, boundary morphology, and tissue heterogeneity, followed by label-preserving appearance augmentation. Because the class map is known before simulation, each synthetic image is paired with a dense mask and metadata without manual annotation. The generated samples therefore provide both image-level diversity and exact pixel-level supervision.

4.2 Experiment 2: Downstream Segmentation Pretraining

The second experiment evaluates whether the generated synthetic data provide transferable supervision for real HFUS segmentation. The real-domain evaluation uses the public Mendeley HFUS dataset, which provides real HFUS images and expert delineations for epidermis and SLEB [7]. Because real labels for dermis, subcutaneous tissue, fascia, and muscle are not available in this dataset, quantitative real-domain evaluation is restricted to epidermis and SLEB. The expanded multilayer label space is therefore demonstrated through the synthetic generation pipeline and dense mask output in Fig. 1, rather than through real-domain multilayer accuracy.

The released SegUNet v0 model [4] is included as a contextual reference and evaluated on the same extracted Mendeley ROI test set. The main comparison focuses on four trainable segmentation architectures: Fresh SegUNet, U-Net [21], DeepLabV3+ [3], and SegFormer [30]. For each trainable architecture, we compare real-only training with synthetic pretraining followed by real fine-tuning. This experiment evaluates the utility of the generated synthetic data, rather than proposing a new segmentation architecture.

All real-domain experiments use 64×128 ROI inputs with the same ROI extraction, split protocol, and Dice/IoU implementation. Performance is measured using class-wise Dice and intersection-over-union (IoU) for epidermis and SLEB, and the main table reports epidermis Dice, SLEB Dice, mean Dice, and mean IoU. All trainable models were optimized using stochastic gradient descent with momentum (SGDM), batch size 8, momentum 0.9, weight decay 5×10^{-4} , and early stopping with patience 40 and minimum delta 10^{-4} . Real-only models were trained for up to 300 epochs, while synthetic-pretraining models used 100 synthetic epochs followed by up to 300 real fine-tuning epochs. Under the model-specific learning-rate policy, Fresh SegUNet used 10^{-3} for synthetic pretraining and 10^{-4} for real training/fine-tuning, whereas U-Net, DeepLabV3+, and SegFormer used 10^{-4} for synthetic pretraining and 10^{-5} for real training/fine-tuning.

5 Results

5.1 Experiment 1: Synthetic Dataset Generation

Fig. 2 shows representative synthetic HFUS images generated under structural and appearance variations. The examples illustrate baseline synthesis, thicker dermis, more irregular boundaries, stronger tissue heterogeneity, and label-preserving rotation augmentation. These samples are intended to show that the

generator can vary both anatomical layout and visual texture while retaining exact pixel-level correspondence with the mask. In particular, the thickness and boundary examples demonstrate controllable structural diversity, whereas the heterogeneity and rotation examples demonstrate appearance diversity used for pretraining.




Fig. 2. Synthetic HFUS examples. From left to right: original synthetic image, thicker dermis variant, irregular boundary variant, stronger tissue heterogeneity, and rotation-based label-preserving augmentation.

5.2 Experiment 2: Downstream Transferability Evaluation

Table 2 reports downstream epidermis/SLEB segmentation on the Mendeley test set. The purpose is to test whether synthetic HFUS images provide transferable supervision, not to establish a new state of the art in segmentation. Synthetic pretraining followed by real fine-tuning achieved performance comparable to real-only training and improved mean Dice/IoU in three of the four trainable architectures. The best mean Dice and mean IoU were obtained by Fresh SegUNet with synthetic pretraining and real fine-tuning, reaching 83.40% and 71.79%. Importantly, synthetic pretraining did not consistently degrade real-domain performance; even when the improvement was small, the pretrained models remained competitive after fine-tuning. Qualitatively, Fig. 3 shows that the SegUNet prediction after synthetic pretraining and real fine-tuning remains similar to the real-only result but yields a slightly more refined boundary in this example. The gains were modest, likely because synthetic images still differ from real scanner-acquired HFUS images in appearance.

6 Discussion

The main contribution of this work is a physics-guided HFUS generation framework for multilayer skin imaging. In contrast to purely data-driven synthesis, it uses explicit acoustic layer maps to generate paired images, dense masks, and metadata.

Table 2. Downstream validation for real-domain epidermis/SLEB segmentation. Released SegUNet v0 is included as a contextual reference; the remaining rows follow the model-specific learning-rate policies described in the text. Mean Dice and mean IoU are arithmetic averages over the epidermis and SLEB classes. Bold indicates the best value for each metric. PT and FT denote pretraining and fine-tuning, respectively.

Method	Training setting	Epi Dice	SLEB Dice	Mean Dice	Mean IoU
Released SegUNet	Released model	78.14%	79.51%	78.82%	65.05%
Fresh SegUNet	Real-only	87.52%	77.75%	82.64%	70.71%
Fresh SegUNet	Synth. PT + real FT	87.91%	78.90%	83.40%	71.79%
U-Net	Real-only	87.61%	76.64%	82.13%	70.04%
U-Net	Synth. PT + real FT	87.79%	78.39%	83.09%	71.35%
DeepLabV3+	Real-only	83.00%	77.56%	80.28%	67.14%
DeepLabV3+	Synth. PT + real FT	82.98%	77.06%	80.02%	66.80%
SegFormer	Real-only	84.67%	80.35%	82.51%	70.28%
SegFormer	Synth. PT + real FT	85.42%	79.91%	82.66%	70.54%




Fig. 3. Representative SegUNet result on a Mendeley case. Top row: original image, expert mask, real-only prediction, and synthetic-pretraining plus real-fine-tuning prediction. Bottom row: original image and corresponding overlays.

The downstream experiment suggests that the generated images provide transferable anatomical and textural priors for real HFUS segmentation. After real fine-tuning, synthetic pretraining achieved performance comparable to real-only training and improved mean Dice/IoU in three of the four trainable architectures, indicating that the images are not merely visually plausible but encode layer-structured cues reusable by real-domain segmentation models.

The moderate improvement indicates that the current simulator captures transferable layer geometry and coarse tissue appearance, but does not yet fully match the HFUS statistics acquired by the scanner. This limitation is not only a matter of sample count: although learned generative models can sample efficiently after training, each physics-based sample requires costly wave-equation simulation. Increasing the number of samples generated by the same simulator may also preserve the same appearance bias. Differences in speckle distribution, attenuation behavior, boundary sharpness, and reflector artifacts may, therefore, explain the architecture-dependent gains.

A key advantage is the generation of dense multilayer labels beyond the real annotations available. Although real-domain evaluation is limited to epidermis and SLEB, the simulator also provides labels for dermis, subcutaneous tissue, fascia, and muscle. Reducing the remaining appearance gap, including overly strong reflector artifacts near the SLEB–dermis transition, is important for extending these initial transfer results toward more reliable multilayer HFUS segmentation.

7 Conclusion

We proposed a physics-guided synthetic HFUS generation framework for skin layer segmentation. The framework uses multilayer acoustic phantoms and k-Wave simulation to generate paired synthetic HFUS images, dense masks, and metadata. It produced structurally diverse synthetic samples and, in downstream segmentation, achieved real-domain performance comparable to real-only training while improving mean Dice/IoU in three of the four trainable architectures. These results suggest that the generated synthetic HFUS images provide transferable cues for real-domain skin layer segmentation. The framework therefore provides a practical basis for studying multilayer HFUS segmentation under controlled synthetic variation while retaining real-domain evaluation for the clinically annotated epidermis/SLEB task. We regard this as an initial validation of the usability of synthetic HFUS rather than a final solution to real-domain segmentation.

Acknowledgments. This work was supported by the Korea Medical Device Development Foundation grant funded by the Korean government (the Ministry of Science and ICT, the Ministry of Trade, Industry and Energy, the Ministry of Health and Welfare, and the Ministry of Food and Drug Safety) (Grant No. RS-2026-25543484).

This research was also supported by the ANCHOR Program through the Gangwon ANCHOR Center, funded by the Ministry of Education (MOE) and Gangwon State (G.S.), Republic of Korea (Grant No. 2026-ANCHOR-10-006).

This research was further supported by the Ministry of Science and ICT (MSIT), Korea, under the National Program in Medical AI Semiconductor (Grant No. 2024-0-00096), supervised by the Institute of Information & Communications Technology Planning & Evaluation (IITP) in 2026.

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

References

1. Barcaui, E.d.O., Carvalho, A.C.P., Piñeiro-Maceira, J., Barcaui, C.B., Moraes, H.: Study of the skin anatomy with high-frequency (22 MHz) ultrasonography and histological correlation. *Radiologia Brasileira* **48**(5), 324–329 (2015). <https://doi.org/10.1590/0100-3984.2014.0028>
2. Chen, A.I., Balter, M.L., Chen, M.I., Gross, D., Alam, S.K., Maguire, T.J., Yarmush, M.L.: Multilayered tissue mimicking skin and vessel phantoms with tunable mechanical, optical, and acoustic properties. *Medical Physics* **43**(6), 3117–3131 (2016). <https://doi.org/10.1118/1.4951729>
3. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder–decoder with atrous separable convolution for semantic image segmentation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018. Lecture Notes in Computer Science*, vol. 11211, pp. 833–851. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01234-2_49
4. Czajkowska, J., Badura, P., Korzekwa, S., Płatkowska-Szczerek, A.: Deep learning approach to skin layers segmentation in inflammatory dermatoses. *Ultrasonics* **114**, 106412 (2021). <https://doi.org/10.1016/j.ultras.2021.106412>
5. Czajkowska, J., Badura, P., Korzekwa, S., Płatkowska-Szczerek, A.: Automated segmentation of epidermis in high-frequency ultrasound of pathological skin using a cascade of DeepLab v3+ networks and fuzzy connectedness. *Computerized Medical Imaging and Graphics* **95**, 102023 (2022). <https://doi.org/10.1016/j.compmedimag.2021.102023>
6. Czajkowska, J., Badura, P., Korzekwa, S., Płatkowska-Szczerek, A., Słowińska, M.: Deep learning-based high-frequency ultrasound skin image classification with multicriteria model evaluation. *Sensors* **21**(17), 5846 (2021). <https://doi.org/10.3390/s21175846>
7. Czajkowska, J., Badura, P., Płatkowska-Szczerek, A., Korzekwa, S.: Data for: Deep learning approach to skin layers segmentation in inflammatory dermatoses. Mendeley Data, Version 1 (2021). <https://doi.org/10.17632/5p7fxjt7vs.1>
8. Czajkowska, J., Juszczak, J., Piejko, L., Glenc-Ambroży, M.: High-frequency ultrasound dataset for deep learning-based image quality assessment. *Sensors* **22**(4), 1478 (2022). <https://doi.org/10.3390/s22041478>
9. Dahan, E., Cohen Indelman, H., Perez-Agosto, A.M., Shiran, C., Avinash, G., Shaked, D., Daniel, N.: CSG: A context-semantic guided diffusion approach in de novo musculoskeletal ultrasound image generation. In: 2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 1–7. IEEE (2025). <https://doi.org/10.1109/EMBC58623.2025.11252664>
10. Frangi, A.F., Tsafaris, S.A., Prince, J.L.: Simulation and synthesis in medical imaging. *IEEE Transactions on Medical Imaging* **37**(3), 673–679 (2018). <https://doi.org/10.1109/TMI.2018.2800298>

11. Freiche, B., El-Khoury, A., Nasiri-Sarvi, A., Hosseini, M.S., Garcia, D., Basarab, A., Boily, M., Rivaz, H.: Ultrasound image generation using latent diffusion models. In: Boehm, C., Mehrmohammadi, M. (eds.) *Medical Imaging 2025: Ultrasonic Imaging and Tomography*. Proceedings of SPIE, vol. 13412, p. 134121F. SPIE (2025). <https://doi.org/10.1117/12.3047788>
12. Jensen, J.A.: Field: A program for simulating ultrasound systems. *Medical & Biological Engineering & Computing* **34**(Suppl. 1), 351–353 (1996)
13. Jensen, J.A., Svendsen, N.B.: Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **39**(2), 262–267 (1992). <https://doi.org/10.1109/58.139123>
14. k-Wave Documentation: Simulating B-mode ultrasound images example. Online documentation, https://www.k-wave.org/documentation/example_us_bmode_linear_transducer.php, last accessed 2026/07/27
15. Lebertre, M., Ossant, F., Vaillant, L., Diridollou, S., Patat, F.: Spatial variation of acoustic parameters in human skin: an in vitro study between 22 and 45 MHz. *Ultrasound in Medicine & Biology* **28**(5), 599–615 (2002). [https://doi.org/10.1016/S0301-5629\(02\)00500-8](https://doi.org/10.1016/S0301-5629(02)00500-8)
16. Levy, J., Barrett, D.L., Harris, N., Jeong, J.J., Yang, X., Chen, S.C.: High-frequency ultrasound in clinical dermatology: a review. *The Ultrasound Journal* **13**, 24 (2021). <https://doi.org/10.1186/s13089-021-00222-w>
17. Loizou, C.P., Pattichis, C.S.: Introduction to ultrasound imaging and speckle noise. In: *Despeckle Filtering Algorithms and Software for Ultrasound Imaging*, pp. 1–20. Springer, Cham (2008). https://doi.org/10.1007/978-3-031-01510-6_1
18. Mast, T.D.: Empirical relationships between acoustic parameters in human soft tissues. *Acoustics Research Letters Online* **1**(2), 37–42 (2000). <https://doi.org/10.1121/1.1336896>
19. Moran, C.M., Bush, N.L., Bamber, J.C.: Ultrasonic propagation properties of excised human skin. *Ultrasound in Medicine & Biology* **21**(9), 1177–1190 (1995). [https://doi.org/10.1016/0301-5629\(95\)00049-6](https://doi.org/10.1016/0301-5629(95)00049-6)
20. Nicolescu, A.C., Ionescu, S., Ancuta, I., Popa, V.T., Lupu, M., Soare, C., Cozma, E.C., Voiculescu, V.M.: Subepidermal low-echogenic band—its utility in clinical practice: A systematic review. *Diagnostics* **13**(5), 970 (2023). <https://doi.org/10.3390/diagnostics13050970>
21. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Lecture Notes in Computer Science, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
22. Shin, M., Peng, Z., Kim, H.J., Yoo, S.S., Yoon, K.: Multivariable-incorporating super-resolution residual network for transcranial focused ultrasound simulation. *Computer Methods and Programs in Biomedicine* **237**, 107591 (2023). <https://doi.org/10.1016/j.cmpb.2023.107591>
23. Shin, M., Seo, M., Yoo, S.S., Yoon, K.: tFUSFormer: Physics-guided super-resolution transformer for simulation of transcranial focused ultrasound propagation in brain stimulation. *IEEE Journal of Biomedical and Health Informatics* **28**(7), 4024–4035 (2024). <https://doi.org/10.1109/JBHI.2024.3389708>
24. Slian, A., Korecka, K., Polańska, A., Czajkowska, J.: Segmentation of skin layers on HFUS images using the attention mechanism. *Computer Methods and Programs in Biomedicine* **263**, 108668 (2025). <https://doi.org/10.1016/j.cmpb.2025.108668>

25. Solano-Cordero, C.M., Encina-Baranda, N., Pérez-Liva, M., Herraiz, J.L.: Recent advances in B-Mode ultrasound simulators. *Applied Sciences* **15**(23), 12535 (2025). <https://doi.org/10.3390/app152312535>
26. Stojanovski, D., Hermida, U., Lamata, P., Beqiri, A., Gomez, A.: Echo from noise: Synthetic ultrasound image generation using diffusion models for real image segmentation. In: Kainz, B., Noble, J.A., Schnabel, J., Khanal, B., Müller, J.P., Day, T. (eds.) *Simplifying Medical Ultrasound*. Lecture Notes in Computer Science, vol. 14337, pp. 34–43. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-44521-7_4
27. Treeby, B.E., Cox, B.T.: k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *Journal of Biomedical Optics* **15**(2), 021314 (2010). <https://doi.org/10.1117/1.3360308>
28. Treeby, B.E., Jaros, J., Rendell, A.P., Cox, B.T.: Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method. *The Journal of the Acoustical Society of America* **131**(6), 4324–4336 (2012). <https://doi.org/10.1121/1.4712021>
29. Vergilio, M.M., Vasques, L.I., Leonardi, G.R.: Characterization of skin aging through high-frequency ultrasound imaging as a technique for evaluating the effectiveness of anti-aging products and procedures: A review. *Skin Research and Technology* **27**(5), 966–973 (2021). <https://doi.org/10.1111/srt.13033>
30. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: SegFormer: Simple and efficient design for semantic segmentation with transformers. In: *Advances in Neural Information Processing Systems*. vol. 34, pp. 12077–12090 (2021)
31. Zama, S., Fujioka, T., Yamaga, E., Kubota, K., Mori, M., Katsuta, L., Yashima, Y., Sato, A., Kawauchi, M., Higuchi, S., Kawanishi, M., Ishiba, T., Oda, G., Nakagawa, T., Tateishi, U.: Clinical utility of breast ultrasound images synthesized by a generative adversarial network. *Medicina* **60**(1), 14 (2024). <https://doi.org/10.3390/medicina60010014>