

# MEDDAP: Medical Dataset Enhancement via Diversified Augmentation Pipeline

Yasamin Medghalchi<sup>1</sup>, Niloufar Zakariaei<sup>1,2</sup>, Arman Rahmim<sup>1,2,3,4</sup>, and Ilker Hacihaliloglu<sup>3,5</sup>

<sup>1</sup> School of Biomedical Engineering, University of British Columbia

<sup>2</sup> Department of Integrative Oncology, BC Cancer Research Institute

<sup>3</sup> Department of Radiology, University of British Columbia

<sup>4</sup> Department of Physics, University of British Columbia

<sup>5</sup> Department of Medicine, University of British Columbia

yasimed@student.ubc.ca, nilouzk@student.ubc.ca, Arman.rahmim@ubc.ca,  
ilker.hacihaliloglu@ubc.ca

**Abstract.** The effectiveness of Deep Neural Networks (DNNs) heavily relies on the abundance and accuracy of available training data. However, collecting and annotating data on a large scale is often both costly and time-intensive, particularly in medical cases where practitioners are already occupied with their duties. Moreover, ensuring that the model remains robust across various scenarios of image capture is crucial in medical domains, especially when dealing with ultrasound images that vary based on the settings of different devices and the manual operation of the transducer. To address this challenge, we introduce a novel pipeline called MEDDAP, which leverages Stable Diffusion (SD) models to augment existing small datasets by automatically generating new informative labeled samples. Pretrained checkpoints for SD are typically based on natural images, and training them for medical images requires significant GPU resources due to their heavy parameters. To overcome this challenge, we introduce USLoRA (Ultrasound Low-Rank Adaptation), a novel fine-tuning method tailored specifically for ultrasound applications. USLoRA allows for selective fine-tuning of weights within SD, requiring fewer than 0.1% of parameters compared to fully fine-tuning only the UNet portion of SD. To enhance dataset diversity, we incorporate different adjectives into the generation process prompts, thereby desensitizing the classifiers to intensity changes across different images. This approach is inspired by clinicians' decision-making processes regarding breast tumors, where tumor shape often plays a more crucial role than intensity. In conclusion, our pipeline not only outperforms classifiers trained on the original dataset but also demonstrates superior performance when encountering unseen datasets.

**Keywords:** Ultrasound · Stable Diffusion · Breast Cancer.

## 1 Introduction

Breast cancer diagnosis heavily relies on mammography as the primary imaging modality due to its widespread availability and established effectiveness in early

detection. However, mammography has inherent limitations, including discomfort for some patients and reduced sensitivity in dense breast tissue [20]. As an alternative, ultrasound has emerged as a valuable supplementary tool in breast imaging. Ultrasound offers several advantages, such as improved comfort during imaging and free from radiation exposure. Moreover, it provides real-time imaging capabilities, making it suitable for further evaluation of suspicious findings detected on mammograms. However, the efficacy of ultrasound is often questioned due to its variable sensitivity and specificity, largely influenced by the operator’s skill and the interpretative complexity of echogenic patterns [27]. In this context, the rise of computational methodologies, particularly deep learning, has been identified as a promising approach to mitigate these challenges. Deep learning’s ability to process and learn from extensive datasets could significantly enhance the diagnostic accuracy of ultrasound imaging in breast cancer [22].

The potential of integrating deep learning into ultrasound imaging is especially significant for improving breast cancer care in resource-limited settings. Ultrasound’s portability and cost-effectiveness, augmented by the advancements in point-of-care ultrasound technology, present a viable alternative in regions where mammography access is constrained [20]. Nevertheless, the application of deep learning in this domain is not without challenges. The accuracy of these models heavily relies on the quality and quantity of the data, where manual collection and the adjustment of image acquisition parameters can introduce significant variability, complicating the development of robust algorithms [28]. Moreover, the scarcity of high-quality, annotated breast ultrasound datasets poses a critical bottleneck, limiting the training and validation of deep learning models [28]. To address these challenges, our research introduces an innovative deep learning (DL) pipeline that leverages stable diffusion (SD) techniques for the augmentation and diversification of ultrasound datasets through prompt engineering [23]. Previously, it was demonstrated that simply expanding the dataset volume is inadequate; a deliberate effort to diversify the dataset is essential to enhance diagnostic precision [8].

Accordingly, our research underscores the vital need for methodically enhancing and diversifying ultrasound datasets, to fully harness deep learning’s diagnostic potential in breast cancer diagnosis. Our specific contributions include: **1-** A novel fine-tuning approach, termed Ultrasound Low-Rank Adaptation (US-LoRA), integrating low-rank adaptation techniques into stable diffusion models, marking the first instance of such methodology within the realm of synthesizing ultrasound data. **2-** A text-based diffusion model guidance aimed at diversifying the generated ultrasound data, particularly focusing on enhancing the diversity of intensity distributions. **3-** Conducting extensive evaluation studies to demonstrate the diversity advantages of our proposed approach in improving classification accuracy.

## 2 Related Works

Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) have been foundational in this space, with VAEs being lauded for their stability and ease of training by optimizing a lower bound on the log-likelihood, while GANs are recognized for their ability to produce high-quality images through adversarial training [21,4]. A combined style augmentation approach that incorporates information from various sources to enhance the training set was proposed in [13]. However, conventional style transfers and traditional augmentation techniques fall short when addressing alterations in shape and they depend on the style image used [29]. Recently, diffusion models have emerged as a potent force for producing synthetic three-dimensional images, presenting notable advancements over conventional generative models such as GANs [7]. Their superior image synthesis capabilities, surpassing those of GANs, stem from their capacity to intricately refine details through a controlled and iterative refinement process. [7,11,18,3,19]. Hence, the utilization of diffusion models in processing ultrasound data has garnered significant interest within the medical imaging community [26,25,17]. These studies collectively illustrate the diverse potential of diffusion models in enhancing ultrasound diagnostics. However, diffusion models often demand considerable computational resources, rendering them less viable in resource-constrained environments—an aspect largely overlooked by prior research in this domain. Furthermore, the exploration of generating diverse diffusion images with varying intensity distribution levels remains unaddressed in the context of ultrasound image analysis.

## 3 Method

### 3.1 USLoRA (Ultrasound Low-Rank Adaptation)

We first explain the use of SD for generating synthetic images, denoted as  $x'_i$ , using specific prompts  $P$  that incorporate the class name  $y'_i$ . In the SD process, an initial image  $x$  undergoes encoding through  $\mathcal{E}$ , yielding  $z = \mathcal{E}\{x\}$ . This is followed by a forward process introducing Gaussian noise  $\epsilon$  at each step  $t$ , resulting in  $Z_t$ . Subsequently,  $Z_T$  and the Text condition  $P$  are combined and input into the denoising U-Net,  $\epsilon_\theta$ , in the reverse process. The SD loss function is the mean squared error (mse), quantifying the difference between the actual noise  $\epsilon$  and the noise predicted by  $\epsilon_\theta(Z_t, t)$ .

$$loss = \mathbb{E}_{\mathcal{E}(x), \epsilon \sim \mathcal{N}(0, I)} [\|\epsilon - \epsilon_\theta(Z_t, t)\|_2^2] \quad (1)$$

The utilization of SD is significantly hindered by its demanding GPU requirements, thus constraining its applicability within the bounds of existing computational resources. To address this challenge, we introduce the USLoRA (Ultrasound Low-Rank Adaptation) method. USLoRA represents the first application of LoRA (Low-Rank Adaptation)[12] for fine-tuning SD specifically for ultrasound images. Conventional fine-tuning approaches necessitate updating the

weight matrix,  $W$ , to align with the pre-trained model’s weight,  $W_0$ . This process, inherently mimics the scale of the pre-trained weight. In standard fine-tuning, the weight matrix adjustment is represented as:

$$W = W_0 + \Delta W, \quad W \in \mathbb{R}^{d \times k}, \quad |W| = |W_0| \quad (2)$$

However, Aghajanyan et al. have demonstrated that the weight change matrices are sparse and possess a low intrinsic rank[1]. This insight gave rise to the LoRA method, which restructures the fine-tuning equation for large language models, like GPT-3, to exploit the low intrinsic rank property[12]. In USLoRA the fine-tuning is performed through the product of two smaller matrices,  $A$  and  $B$ , with rank  $r$ :

$$W = W_0 + AB, \quad A \in \mathbb{R}^{d \times r}, \quad B \in \mathbb{R}^{r \times k} \quad r \ll d \quad (3)$$

In USLoRA, for fine-tuning SD for Ultrasound image generation, we harness the unique capabilities of LoRA [12] specifically within the cross-attention layers of UNet architecture. The decision to freeze the remaining model components is strategic, focusing computational efforts precisely where they are most impactful - enhancing image representations guided by input prompts. These large models are typically trained on extensive datasets and demonstrate some capability in comprehending both images and prompts. Our objective is just to establish a connection between the provided prompts and images within our ultrasound dataset, necessitating fine-tuning of the cross-attention layers using LoRA for weight matrices  $W_q, W_k, W_v, W_o$ . Finally, we fine-tune SD with less than 0.1% parameters compared to fine-tuning the whole UNet module of SD.

### 3.2 Prompt Engineering

We utilize SD as a text-to-image generator, requiring textual descriptions as input. To enrich our dataset effectively, we enhance input prompts with descriptive adjectives. Our adjective selection is driven by two observations. First, we consider the variability of ultrasound image intensity arising from device differences and transducer orientation. Second, we aim to challenge classifiers with diverse clinical scenarios while recognizing that tumor detection relies more on structural features than intensity. Given the limited dataset size, we ensure these additions remain general rather than overly specific to avoid overwhelming the model with excessive detail. Our augmentation strategy includes adjectives such as 'colorful,' 'solarized,' and 'stylized'. Our prompt structure follows the format ' {adj} ultrasound image of benign/malignant/no tumor in the breast'. By integrating class labels (benign/malignant) into the SD model, we mitigate errors stemming from ambiguity in the input text. Moreover, the inclusion of adjectives enhances the diversification of SD model generations, capturing a wide array of discriminative features that can enhance classification accuracy.

## 4 Experiments and Results

**Dataset.** In this study, we utilize three open-source breast ultrasound datasets. Dataset A was exclusively employed for training the SD model and the classifiers utilized to evaluate the enhancements attained with our augmented data for the downstream task. Dataset B was only utilized as test data for the downstream task of classification.

**Dataset A.** Comprising images from female patients aged 25 to 75, aimed at detecting breast cancer [2], this dataset, collected in 2018, includes 780 ultrasound images from 600 patients, with an average size of 500x500 pixels. Categorized into normal, benign, and malignant classes, it consists of 133, 437, and 210 images, respectively. We partition the dataset into training (545 images), validation sets (115 images), and testing (120 images) at the patient level [2]. The training set is utilized for training and fine-tuning SD. Given the limited dataset size, we employ the five-fold cross-validation technique to conduct our experiments.

**Dataset B.** Dataset B, an open-access dataset available on Kaggle[16], includes ultrasound images depicting benign and malignant breast cancers. These images have undergone augmentation via rotation and sharpening to augment the dataset size. For testing the trained classifiers, we utilize the test set from this dataset, comprising 500 images for benign cases and 400 images for malignant cases [6].

**Implementation Details.** We based our experiments on a pre-trained model[5], initialized with Stable-Diffusion-v-1-2 checkpoint weights and fine-tuned for 225k steps at 512x512 resolution on "laion-aesthetics v2 5+". We conducted experiments to fine-tune the proposed USLoRA method using hyperparameters of 2, 4, and 8. Initially, we generated images equivalent in size to the training set and compared them based on the FID score [24]. The resulting FID scores were 0.463, 0.357, and 0.513 for ranks of 2, 4, and 8, respectively. Consequently, we determined that a rank of 4 yields the optimal outcome. Training USLoRA on SD for Dataset A is done with 100 epochs of AdamW optimizer with a learning rate of  $10^{-4}$ , and the batch size is 1, requiring 6.9 GB of GPU memory and less than 5 hours of one NVIDIA GeForce RTX 2080 Ti GPU. The input image dimensions were set to  $224 \times 224$ . To demonstrate the utility of synthesized ultrasound images in enhancing the performance of a supervised downstream classification task, we leverage ResNet34 [10], SqueezeNet1.1 [15], and DenseNet121 [14] architectures, widely recognized as baseline models in medical imaging. The classifiers were trained for 100 epochs with a batch size of 32. In our experimental setup, we employed the Adam optimizer for its efficiency in handling sparse gradients and its adaptability to large-scale data problems. The learning rate was set to 0.001, following conventional best practices for achieving a balance between convergence speed and training stability. It's important to note that while these architectures are integral to our evaluation, they do not represent the central focus of our study. Instead, our primary contribution lies in the synthesis of images, which can augment a range of CNN-based network architectures.

**Effect of adjectives in classification performance.** In this section, we commence by training the proposed text-guided SD model on the train set of Dataset A. Notably, the generation process is confined to the train set of Dataset A exclusively, with no utilization of images from the validation or test set of Dataset A during SD model training.

Table 1: Analyzing the impact of various adjectives with different extension ratios on classifier accuracy. "" means SD was trained without an adjective. % indicates the size of the synthesized data included in the training data. Test data comprised real ultrasound data from Dataset A. A lower FID indicates better-quality images. **Bold text indicates the best results achieved.**

Adjective	FID	Densenet121			ResNet34			Squeezenet1.1		
		50%	100%	200%	50%	100%	200%	50%	100%	200%
""	0.357	86.66	80.83	80.83	85	81.66	83.33	78.33	69.17	75
"Colorful"	0.374	86.66	86.66	83.33	85	81.66	79.17	80	68.33	71.67
"Stylized"	0.367	83.33	87.5	81.67	<b>87.5</b>	80	81.67	76.67	78.33	75
"High-contrast"	0.561	<b>90</b>	83.33	83.33	85	76.66	83.33	<b>82.5</b>	81.67	70.83
"Low-contrast"	0.503	84.16	82.5	80.83	85.83	79.16	77.5	67.5	72.5	77.5
"Posterized"	0.435	80.83	87.5	86.67	85.83	80	79.17	69.17	78.33	74.17
"Sheared"	0.499	83.33	87.5	84.17	83.33	85	76.67	75.83	67.5	70.83
"Solarized"	<b>0.26</b>	86.66	85	83.33	<b>87.5</b>	83.33	84.17	74.17	77.5	73.33
"Bright"	0.663	84.16	83.33	84.17	84.16	75	77.5	74.17	73.33	73.33
"Dark"	0.354	85.83	87.5	82.5	76.66	78.33	78.33	77.5	74.17	72.5

Table 2: Comparison of classification performance between classifiers trained solely on real data and those trained on a combination of real and synthetic data. Test data comprised real ultrasound data from Dataset A (paired t-test <0.05). **Bold text indicates the best results achieved.**

Train Dataset		Accuracy	Sensitivity	Specificity	Precision	F1
DenseNet121	Real-ultrasound only	87.5	0.87	0.92	0.88	0.87
	Real+50%-High-contrast	<b>90</b>	<b>0.90</b>	<b>0.94</b>	<b>0.89</b>	<b>0.89</b>
ResNet34	Real-ultrasound only	79.17	0.77	0.87	0.81	0.79
	Real+50%-Stylized	<b>87.5</b>	<b>0.88</b>	<b>0.93</b>	<b>0.85</b>	<b>0.86</b>
Squeezenet1.1	Real-ultrasound only	72.5	0.68	0.86	0.67	0.67
	Real+50%-High-contrast	<b>82.5</b>	<b>0.80</b>	<b>0.89</b>	<b>0.82</b>	<b>0.81</b>

As a subsequent procedure, we generate images in three varying sizes: 50%, 100%, and 200% of the Dataset A test data size(545), accompanied by nine adjectives and a scenario devoid of any adjective. These images are subsequently integrated with the real image train data from Dataset A to establish a new training dataset,

which is employed for training the classifiers. The validation and test datasets consistently comprise real ultrasound images from Dataset A, as elucidated in the 'Dataset' section. In the first row of Table 1, the absence of an entry ("") indicates that no adjective was utilized during training. As demonstrated in Table 1, where we display the average accuracy, the utilization of adjectives consistently outperforms scenarios without adjectives in the majority of cases. Table 2 presents quantitative results comparing the performance of classifiers when using real ultrasound data only versus the best synthetic datasets generated from the optimal performance of each classifier as shown in Table 1. Our proposed pipeline demonstrates superior performance over the original datasets in terms of accuracy, sensitivity, specificity, precision, and F1 score (Table 2) (paired t-test  $<0.05$ ).

**Selection of adjectives.** In our pipeline, a common question pertains to the selection of adjectives. While there is inherent randomness in the generation process, we sought to elucidate the selection of adjectives through experiments.

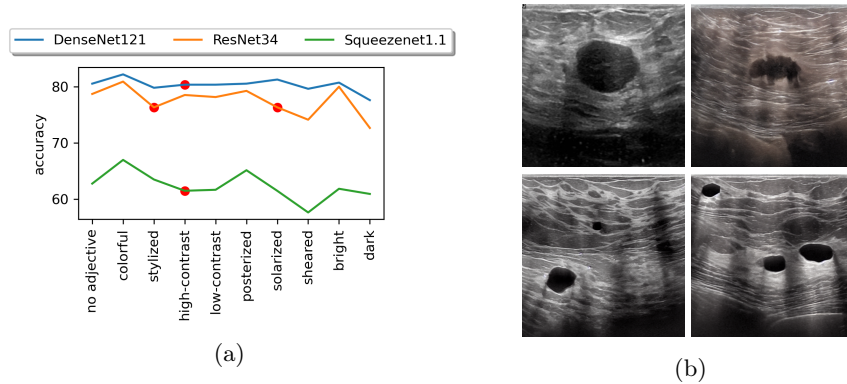


Fig. 1: (a) Accuracy results when networks are trained on real data only and tested on synthetic datasets only. (b) Some examples of synthesized images. Top-left: Original B-mode ultrasound data, synthesized images using top-right: Colorful, bottom-left: Solarized, bottom-right: High-contrast adjectives.

By testing synthetic datasets with various adjectives on classifiers trained on the real ultrasound train set only, we aimed to shed light on this aspect. Figure 1-(a) illustrates the outcomes, revealing that the best adjectives (depicted by red circles) do not consistently yield the highest or lowest performance. Notably, achieving the best accuracy implies that the generated images closely resemble the real ultrasound data, potentially indicating that the SD model is lacking new information (diversity). Conversely, the worst performance manifests that the generated images are distant from the distribution of the real ultrasound data. In contrast, we aim to guide existing diffusion models towards a broader comprehension of classes by emphasizing the distinguishing characteristics neces-

sary for discrimination, all while harnessing the full diversity of the underlying generative model. Figure 1-(b) shows qualitative results obtained using a sample of adjectives.

**Effect of our proposed method in robustness.** In this section, we assess the efficacy of our method in accurately classifying unseen datasets, specifically Dataset B [6]. Table 3 illustrates that our top-performing adjectives, identified from Table 1 for each classifier, consistently outperform the classifiers trained with real data only across all scenarios.

Table 3: Accuracy results evaluating the robustness of our pipeline on unseen test Dataset B (paired t-test  $<0.05$ ). **Bold text indicates the best results achieved.**

	Train Dataset	Adjective	Test Set B Accuracy
DenseNet121	Real-ultrasound only	-	79.66
	Real+50% Augmented	"High-Contrast"	<b>83.22</b>
ResNet34	Real-ultrasound only	-	77.44
	Real+50% Augmented	"Stylized"	<b>80.44</b>
Squeezenet1.1	Real-ultrasound only	-	64.44
	Real+50% Augmented	"High-Contrast"	<b>72.89</b>




Fig. 2: Heat map results depict classification network attention. Columns from right to left: Real ultrasound data, heat map output when classification network (Resnet34 top, DenseNet121 bottom) was trained solely on real ultrasound dataset, and heat map output when classification network (Resnet34 top, DenseNet121 bottom) was trained on real ultrasound dataset augmented with 50% "Stylized" adjective (top) and 50% "High-Contrast" adjective (bottom).



Analyzing the heat-map images depicted in Figure 2, we note that models trained on original datasets primarily focus on textual elements (part a) and the ultrasound image’s shadow (part b) rather than tumors. The classifier’s focus on the text exemplifies a typographic attack [9], where the classifier prioritizes textual elements over the depiction of tumors in the ultrasound image. In contrast, we observe that integrating synthetic datasets with real data enhances the reliability of the decision-making process.

## 5 Conclusion

This study proposes a novel pipeline for expanding small-scale ultrasound datasets by incorporating informative images rather than simply increasing dataset size. This approach not only enhances classification accuracy but also bolsters the network’s robustness and reliability by diversifying the dataset. Our proposed method addresses the challenges associated with training diffusion models, which typically demand substantial memory resources, thereby facilitating broader access for fine-tuning this extensive model to address various problems.

## References

1. Aghajanyan, A., Zettlemoyer, L., Gupta, S.: Intrinsic dimensionality explains the effectiveness of language model fine-tuning. arXiv:2012.13255 [cs] (2020)
2. Al-Dhabyani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. Data in Brief (2020)
3. Ali, H., Murad, S., Shah, Z.: Spot the fake lungs: Generating synthetic medical images using neural diffusion models. In: Irish Conference on Artificial Intelligence and Cognitive Science. pp. 32–39. Springer (2022)
4. Alsinan, A.Z., Rule, C., Vives, M., Patel, V.M., Hacihaliloglu, I.: Gan-based realistic bone ultrasound image and label synthesis for improved segmentation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23. pp. 795–804. Springer (2020)
5. CompVis: Stable Diffusion V1.4 - Original (2023), <https://huggingface.co/CompVis/stable-diffusion-v-1-4-original>, accessed: [Insert Date of Access]
6. Ultrasound breast images for breast cancer, <https://www.kaggle.com/datasets/vuppalaadithyasairam/ultrasound-breast-images-for-breast-cancer/data>
7. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
8. Fujioka, T., Mori, M., Kubota, K., Oyama, J., Yamaga, E., Yashima, Y., Katsuta, L., Nomura, K., Nara, M., Oda, G., et al.: The utility of deep learning in breast ultrasonic imaging: a review. *Diagnostics* **10**(12), 1055 (2020)
9. Goh, G., Cammarata, N., Voss, C., Carter, S., Petrov, M., Schubert, L., Radford, A., Olah, C.: Multimodal neurons in artificial neural networks. *Distill* **6**(3), e30 (2021)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)

11. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
12. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. *arXiv:2106.09685v2 [cs.CL]* (2021)
13. Huang, B., Xu, Z., Chan, S.C., Liu, Z., Wen, H., Hou, C., Huang, Q., Jiang, M., Dong, C., Zeng, J., et al.: A style transfer-based augmentation framework for improving segmentation and classification performance across different sources in ultrasound images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 44–53. Springer (2023)
14. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017)
15. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size. *arXiv:1602.07360* (2016)
16. Kaggle, <https://www.kaggle.com/docs/datasets>
17. Katakis, S., Barotsis, N., Kakotartitis, A., Tsiganos, P., Economou, G., Panagiotopoulos, E., Panayiotakis, G.: Generation of musculoskeletal ultrasound images with diffusion models. *BioMedInformatics* **3**(2), 405–421 (2023)
18. Kazerouni, A., Aghdam, E.K., Heidari, M., Azad, R., Fayyaz, M., Hacıhaliloglu, I., Merhof, D.: Diffusion models for medical image analysis: A comprehensive survey. *arXiv preprint arXiv:2211.07804* (2022)
19. Khader, F., Mueller-Franzes, G., Arasteh, S.T., Han, T., Haarburger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baessler, B., Foersch, S., et al.: Medical diffusion–denoising diffusion probabilistic models for 3d medical image generation. *arXiv preprint arXiv:2211.03364* (2022)
20. Lee, C.I., Lehman, C.D.: Digital breast tomosynthesis and the challenges of implementing an emerging breast cancer screening technology into clinical practice. *Journal of the American College of Radiology* **13**(11), R61–R66 (2016)
21. Liang, J., Yang, X., Huang, Y., Li, H., He, S., Hu, X., Chen, Z., Xue, W., Cheng, J., Ni, D.: Sketch guided and progressive growing gan for realistic and editable ultrasound image synthesis. *Medical Image Analysis* **79**, 102461 (2022)
22. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. *Medical image analysis* **42**, 60–88 (2017)
23. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2021)
24. Seitzer, M.: pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid> (August 2020), version 0.3.0
25. Stevens, T.S., Meral, F.C., Yu, J., Apostolakis, I.Z., Robert, J.L., Van Sloun, R.J.: Dehazing ultrasound using diffusion models. *IEEE Transactions on Medical Imaging* (2024)
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. *arXiv preprint arXiv:2305.05424* (2023)
27. Svensson, W.E., Stewart, V.R.: Advanced applications of breast ultrasound. *Breast Cancer* pp. 70–98 (2010)
28. Yap, M.H., Pons, G., Marti, J., Ganau, S., Sentis, M., Zwiggelaar, R., Davison, A.K., Marti, R.: Automated breast ultrasound lesions detection using convolutional neural

- networks. *IEEE journal of biomedical and health informatics* **22**(4), 1218–1226 (2017)
29. Zhao, B., Yu, S., Ma, W., Yu, M., Mei, S., Wang, A., He, J., Yuille, A., Kortylewski, A.: Ood-cv: a benchmark for robustness to out-of-distribution shifts of individual nuisances in natural images. In: *European Conference on Computer Vision*. pp. 163–180. Springer (2022)