

# From Majority to Minority: A Diffusion-based Augmentation for Underrepresented Groups in Skin Lesion Analysis

Janet Wang, Yunsung Chung, Zhengming Ding, and Jihun Hamm

Tulane University

{swang47, ychung3, zding1, jhamm3}@tulane.edu

**Abstract.** AI-based diagnoses have demonstrated dermatologist-level performance in classifying skin cancer. However, such systems are prone to under-performing when tested on data from minority groups that lack sufficient representation in the training sets. Although data collection and annotation offer the best means for promoting minority groups, these processes are costly and time-consuming. Prior works have suggested that data from majority groups may serve as a valuable information source to supplement the training of diagnostic tools for minority groups. In this work, we propose an effective diffusion-based augmentation framework that maximizes the use of rich information from majority groups to benefit minority groups. Using groups with different skin types as a case study, our results show that the proposed framework can generate synthetic images that improve diagnostic results for the minority groups, even when there is little or no reference data from these target groups. The practical value of our work is evident in medical imaging analysis, where under-diagnosis persists as a problem for certain groups due to insufficient representation. Our implementation detail is available at <https://github.com/janet-sw/skin-diff>.

**Keywords:** Skin Lesion Analysis · Diffusion Models · Data Augmentation.

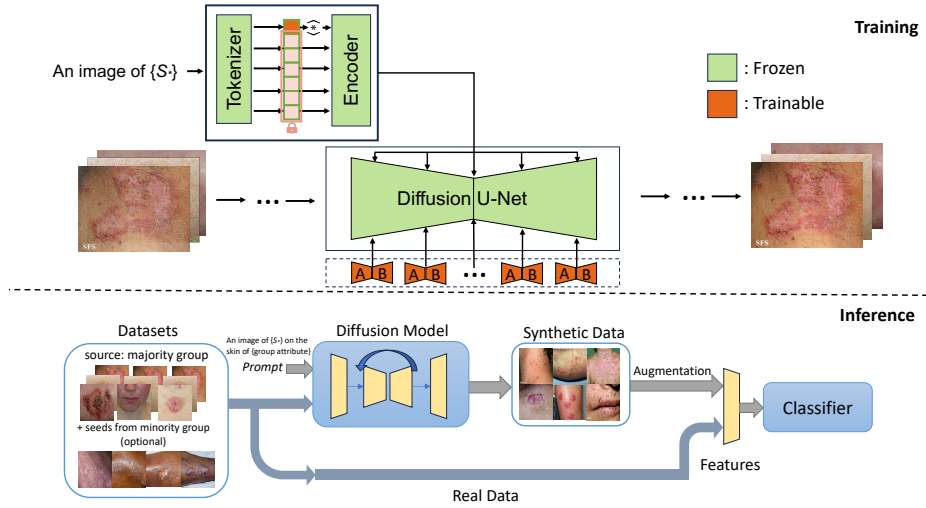
## 1 Introduction

AI-assisted diagnostic systems demonstrate expert-level capability in classifying skin cancers, often identified visually [6,16,2]. These systems can potentially contribute to teledermatology as diagnostic and decision-support tools, enhancing diagnostic accessibility in rural areas [3]. However, despite such success, recent studies have highlighted their susceptibility to under-diagnosing minority groups, such as those with underrepresented skin types, hindering their ability to generalize across different demographic groups [4,11]. Although the majority group contains rich lesion information, directly training models for cross-color classification using this data is challenging due to the domain gap caused by varying skin types [24]. Prior research has suggested using synthetic images generated from majority groups to supplement the training of AI for minority groups [19].

Augmenting skin condition data with synthetic images has been explored, owing to its potential to address common challenges for skin lesion analysis, such as data privacy, imbalance, and scarcity. Notably, Generative Adversarial Networks (GANs) [10] and Diffusion Models (DMs) [5] have emerged as leading techniques for generating high-quality skin lesion images. While GANs have successfully produced photorealistic synthetic images, their generation is uncontrollable [9,18]. On the other hand, DMs pre-trained on extensive web data have enabled higher controllability over image generation through the guidance of textual prompts, allowing for the creation of diverse and high-fidelity images of target skin conditions and skin types.

Existing studies have tried diffusion models to augment minority skin types using two public datasets: Diverse Dermatology Images (DDI) [4] and Fitzpatrick17k [11]. Each image in these datasets is annotated with skin type labels based on the Fitzpatrick scoring system [7]. In their work, [21] generated multiple synthetic images for each real image using Stable Diffusion [20] and then trained the classifier on a dataset including real and synthetic data. They found that diffusion models can enhance diagnosis accuracy across skin types in binary malignancy classification on the DDI dataset, though the number of real images is the key driver in performance. Additionally, [22] sampled a small number of seed images with skin types at the ends of the Fitzpatrick spectrum (FST I-II and FST V-VI) and carefully cropped the disease pathology in them, before generating synthetic data from the seeds using OpenAI DALL·E 2’s inpainting feature. They conducted class-wise data augmentation by incorporating synthetic images of the target condition and minority skin type into the real training set. Other related studies have focused on internal datasets [1,15].

Despite these advancements, the potential to leverage diffusion models’ knowledge about skin variation and the rich lesion information from majority groups to benefit minority groups remains underexplored. In this work, we propose a novel diffusion-based augmentation framework capable of learning skin lesion concepts from majority groups and generating images to improve classification performance for minority groups. Unlike current works that assume the existence of data from minority groups, we hypothesize that the information gained from majority groups and the diffusion model’s pre-trained knowledge is sufficient to generate useful synthetic data. We test our hypothesis in a challenging multi-condition classification task. The framework is illustrated in Fig. 1. We conduct our experiments on the Fitzpatrick17k dataset, which includes lesions that are less familiar to diffusion models than common skin cancer. This dataset has a skewed skin type distribution, with light skin types (FST I-II) being significantly more than dark skin types (FST V-VI), thus forming majority and minority groups. Our investigation focuses on images from both groups and is structured around three scenarios with increasing difficulty: **(i)** the training source includes some data from both groups; **(ii)** there is limited data from the minority group in the training source; and **(iii)** the training source lacks data from the minority group. Through extensive experiments and analysis, we found that:



**Fig. 1.** Overview of the proposed augmentation framework. The framework pairs each training image with a textual prompt describing the condition as an input to train a latent diffusion model. Embeddings associated with new lesion concepts are found through Textual Inversion. Compact matrices  $A$  and  $B$  are optimized via LoRA to facilitate training with the new embeddings. During inference, the trained model produces synthetic images from the training set that mainly features the majority groups via image-to-image generation, thus conditioned on visual cues of lesions from images and textual prompts describing the target condition and group attributes.

- Our proposed method effectively leverages lesion information from the majority group to generate synthetic images that can improve classification for the minority group across all settings, even without reference data from the minority group.
- Using synthetic images generated by our method to train classifiers consistently outperforms training with real images across various architectures. Further improvement is observed when combining real and synthetic data.
- Our method is sensitive to information from the minority group. A notable improvement can be observed when even a few examples from the minority group are added to the training set.

## 2 Methods

In this section, we will introduce key techniques that have been adapted for skin disease datasets in our proposed augmentation framework.

**Latent Diffusion Models** We implement our method using Latent Diffusion Models (LDMs) [20], a class of Denoising Diffusion Probabilistic Models (DDPMs) [13] that operate in the latent space of an autoencoder, to enable

<i>FST</i>	Basal Cell Carcinoma	Folliculitis	Nematode Infection	Neutrophilic Dermatoses	Prurigo Nodularis	Psoriasis	Squamous Cell Carcinoma	<i>Total</i>
<i>I</i>	85	30	15	70	7	113	100	420
<i>II</i>	156	97	56	115	28	232	180	864
<i>V</i>	24	31	32	31	29	64	40	251
<i>VI</i>	7	9	12	15	9	21	23	96
Total	272	167	115	231	73	430	343	<b>1631</b>

**Table 1.** Sample distribution across skin conditions by Fitzpatrick Skin Type.

DDPM training with limited computational resources. LDMs include two core components: a pre-trained autoencoder and a diffusion model. In our study, the encoder of the autoencoder  $\mathcal{E}$  encodes skin lesion images  $x \in \mathcal{D}_x$  into a latent representation  $z = \mathcal{E}(x)$ , while the decoder  $D$  maps the latent representations back to images, such that  $D(\mathcal{E}(x)) \approx x$ . The diffusion model is trained to generate representations conditioned on prompts describing skin disease and skin type, within the learned latent space. Let  $c_\theta(y)$  be a model that maps a conditioning input  $y$  into a vector. We then learn the conditional LDM via

$$L_{LDM} := \mathbb{E}_{z \in \mathcal{E}(x), y, \epsilon \in \mathcal{N}(0,1), t} \left[ \|\epsilon - \epsilon_\theta(z_t, t, c_\theta(y))\|_2^2 \right], \quad (1)$$

where  $t$  is the time step,  $z_t$  is the latent noise at time  $t$ ,  $\epsilon$  is the unscaled noise sample, and  $\epsilon_\theta$  is the denoising network.

**Concept Discovery via Textual Inversion** Our proposed framework leverages Textual Inversion [8] to capture a unique embedding that accurately represents the targeted skin lesion concept from training data. Skin lesion images paired with a string containing a placeholder word (e.g., ‘An image of  $\{S_*\}$ ’) are used to guide the learning of a new lesion embedding for the generative model. In particular, the optimal embedding  $v_*$  that encapsulates the lesion concept  $S_*$  is derived by minimizing the reconstruction loss,

$$v_* = \operatorname{argmin}_v \mathbb{E}_{z \sim \mathcal{E}(x), y, \epsilon \sim \mathcal{N}(0,1), t} \left[ \|\epsilon - \epsilon_\theta(z_t, t, c_\theta(y, S_*))\|_2^2 \right], \quad (2)$$

where the same training scheme as the original LDM model is used, with  $c_\theta$  and  $\epsilon_\theta$  fixed.

**Fine-grained Detail Enhancement with LoRA** To enhance efficiency in fine-tuning LDM, we employ Low-Rank Adaptation (LoRA) [14] in our framework, with the discovered tokens after textual inversion. This fine-tuning strategy freezes the pre-trained model weights and introduces two compact matrices  $A$  and  $B$ , where  $A \in \mathbb{R}^{n \times r}$ ,  $B \in \mathbb{R}^{r \times n}$ . The adaptation matrices  $AB$  are integrated into the attention layers to capture fine visual details of the skin lesion that were not initially present in the pre-trained model, with target embedding  $v_*$ . The optimization is formulated as

$$L := \mathbb{E}_{z \sim \mathcal{E}(x), y, \epsilon \sim \mathcal{N}(0,1), t} \left[ \|\epsilon - \epsilon_{\theta_{AB}}(z_t, t, c_{\theta_{AB}}(y, v_*))\|_2^2 \right]. \quad (3)$$

Architecture	Train Type	Train Size	Accuracy	Precision	Recall	F1 Score
VGG-16	real	1519	$70.24 \pm 0.12$	$72.49 \pm 0.37$	$69.58 \pm 0.30$	$70.48 \pm 0.13$
	syn	1519 * 5	$75.00 \pm 0.64$	$75.77 \pm 0.39$	$73.17 \pm 0.29$	$72.42 \pm 0.61$
	real + syn	1519 * 6	$77.98 \pm 0.40$	$81.51 \pm 0.25$	$78.87 \pm 0.23$	$77.45 \pm 0.29$
ResNet-18	real	1519	$68.45 \pm 0.42$	$69.42 \pm 0.57$	$69.05 \pm 0.42$	$67.94 \pm 0.35$
	syn	1519 * 5	$69.05 \pm 0.36$	$69.57 \pm 0.82$	$69.05 \pm 0.42$	$68.02 \pm 0.33$
	real + syn	1519 * 6	$71.36 \pm 0.69$	$69.02 \pm 0.48$	$68.45 \pm 0.40$	$67.56 \pm 0.59$
ViT-B-16	real	1519	$70.38 \pm 0.42$	$73.72 \pm 0.69$	$70.82 \pm 0.39$	$70.61 \pm 0.59$
	syn	1519 * 5	$74.19 \pm 0.37$	$77.89 \pm 1.01$	$74.04 \pm 0.85$	$73.58 \pm 0.53$
	real + syn	1519 * 6	$78.65 \pm 0.53$	$81.57 \pm 0.47$	$79.17 \pm 0.84$	$78.24 \pm 0.64$

**Table 2.** This table presents the results when the training set includes non-flexible images from the minority group (291 of FST V-VI) and the majority group (1228 FST I-II). The test set is a flexible subset of the minority group (56 of FST V-VI), uniformly distributed across the 7 conditions. Here, “real” indicates that the classifier is trained solely on real images, while “syn” means that it is trained exclusively on synthetic images generated by our framework. Accordingly, “real+syn” means the subsequent classifier is trained on a combination of both.

### 3 Experiments

We conduct our experiments using the Fitzpatrick17k dataset, where each image is annotated with a condition and a Fitzpatrick Skin Type (FST) label. In line with [22], our analysis narrows down to a subset of the Fitzpatrick17k dataset, encompassing 7 conditions (Table 1). These conditions were selected because they represent the largest sample sizes at the ends of the Fitzpatrick Skin Type (FST I-II or V-VI) spectrum. Unlike [22], our study excludes intermediate skin types (FST III-IV), to explore the efficacy of our diffusion-based augmentation in a more challenging and explainable way. We randomly sample 8 images for each condition from the lightest (FST I-II) and darkest (FST V-VI) skin type groups, resulting in a **flexible subset** of 56 images for each group.

We examine three scenarios: **(i)** the training set includes images of dark and light skin types, and the test set features the uniformly distributed flexible subset across the 7 conditions; **(ii)** the training set predominantly includes light-skinned images and a few dark-skinned images, while the test set consists of dark-skin data; **(iii)** the training set lacks dark-skinned images entirely, while the test set comprises dark-skinned images. In all scenarios, we generate 5 synthetic images for each real one in the training set during inference, using the fine-tuned model, as illustrated in Fig. 1. In scenario **(i)**, to ensure a sufficient number of examples for both majority and minority groups in the training set, we designate the flexible subset of dark skin as the test set and use remaining non-flexible images for generator and classifier training. This setting also serves as the basis for hyperparameter tuning of the diffusion model, with the selected hyperparameters being fixed for subsequent experiments.

**Implementation Details** In each setting, we randomly sampled 5 flexible subsets and repeated the experiment 5 times. We used the Stable Diffusion 2.1 base [20] and the Diffusers library [17] for fine-tuning the diffusion model and

Architecture	Train Type	Train Size	Accuracy	Precision	Recall	F1 Score
VGG-16	real	1284	$58.79 \pm 0.10$	$58.90 \pm 0.03$	$58.26 \pm 0.05$	$56.98 \pm 0.04$
	syn	1284 * 5	$62.86 \pm 0.15$	$61.49 \pm 0.15$	$63.32 \pm 0.13$	$61.57 \pm 0.24$
	real + syn	1284 * 6	$63.66 \pm 0.11$	$62.80 \pm 0.12$	$64.08 \pm 0.08$	$62.72 \pm 0.18$
ResNet-18	real	1284	$50.31 \pm 0.30$	$50.45 \pm 0.44$	$51.27 \pm 0.30$	$48.23 \pm 0.32$
	syn	1284 * 5	$56.36 \pm 0.16$	$56.58 \pm 0.13$	$59.78 \pm 0.08$	$55.58 \pm 0.10$
	real + syn	1284 * 6	$61.33 \pm 0.17$	$59.75 \pm 0.12$	$62.52 \pm 0.17$	$59.94 \pm 0.17$
ViT-B-16	real	1284	$62.03 \pm 0.19$	$63.09 \pm 0.13$	$62.02 \pm 0.04$	$61.05 \pm 0.04$
	syn	1284 * 5	$68.83 \pm 0.19$	$70.07 \pm 0.03$	$68.22 \pm 0.25$	$68.34 \pm 0.19$
	real + syn	1284 * 6	$71.20 \pm 0.07$	$71.61 \pm 0.19$	$71.53 \pm 0.01$	$71.17 \pm 0.13$

**Table 3.** Classification results when the training set contains a few reference images from the dark-skinned flexible subset (56 of FST V-VI) and non-flexible light-skinned images (1228 of FST I-II). The test set includes all other dark-skinned images (291 of FST V-VI) outside the flexible subset.

generating synthetic images. For classifier backbones, we utilized pre-trained VGG-16 [23], ResNet-18 [12], and ViT-B-16 [25] and trained each classifier using the Adam optimizer with an initial learning rate of 1e-3 and transformations as in [11]. A weight-based sampler and StepLR scheduler were applied. All experiments were conducted on two NVIDIA GeForce RTX 3090s.

## 4 Results

To assess the efficacy of our augmentation framework across the three settings, we train the classifier on data that includes real images only, synthetic images only, or a combination of both, respectively. Our evaluation is based on four metrics: accuracy, precision, recall, and F1. First, in the setting with some images from both majority and minority groups in the training set, we observe that synthetic data enhances performance across all architectures (Table 2). Specifically, classifiers trained on synthetic images consistently outperform those trained solely on real ones, and the combination of both types of data for training yields further improvements.

This trend of consistent improvement is also evident in the more challenging scenarios where there are little or no reference images from the minority groups (Tables 3 and 4) in the training set. Notably, significant improvement is observed when just a few reference images from the minority group are available in the training set for image generation and classification. The transformer-based classifier demonstrates a larger improvement gap over the real image baseline than the CNN-based models. In the most challenging setting, with no reference images from the minority group, the improvement margin narrowed, suggesting that our pipeline effectively maximizes the use of limited information from the flexible subset of the minority group during training. Despite these challenges, the sustained improvements in the third setting validate our framework’s effectiveness in transferring information across groups. Examples of real and synthetic image pairs for each condition are presented in Fig. 2. Qualitatively, the synthetic

Architecture	Train Type	Train Size	Accuracy	Precision	Recall	F1 Score
VGG-16	real	1228	55.58 ± 0.10	54.60 ± 0.11	51.84 ± 0.29	51.97 ± 0.45
	syn	1228 * 5	57.62 ± 0.09	56.62 ± 0.15	55.42 ± 0.20	55.36 ± 0.22
	real + syn	1228 * 6	58.08 ± 0.08	57.36 ± 0.13	55.78 ± 0.17	55.85 ± 0.23
ResNet-18	real	1228	49.42 ± 0.36	49.79 ± 0.44	48.30 ± 0.23	47.42 ± 0.30
	syn	1228 * 5	53.47 ± 0.32	52.39 ± 0.30	53.79 ± 0.30	51.97 ± 0.29
	real + syn	1228 * 6	56.50 ± 0.15	55.72 ± 0.15	55.10 ± 0.16	54.76 ± 0.18
ViT-B-16	real	1228	57.66 ± 0.35	61.89 ± 0.68	55.45 ± 0.36	55.96 ± 0.51
	syn	1228 * 5	58.94 ± 0.23	62.79 ± 0.20	55.62 ± 0.22	56.53 ± 0.18
	real + syn	1228 * 6	60.96 ± 0.31	64.57 ± 0.01	57.51 ± 0.29	58.92 ± 0.31


**Table 4.** Classification results when no image from the minority group is in the training set, which only has non-flexible images of the majority group (1228 of FST I-II). The test set has non-flexible images of the minority group (291 of FST V-VI).

images generated by our augmentation framework introduce more diversity to the training sets, including variations in skin color and lesion patterns.

To further investigate our framework’s generation capabilities, we conduct an ablation study comparing our framework with various generation strategies. This study focuses on the first setting, where the test set consists of light- or dark-skinned flexible images (56 for each type). We first examine the Stable Diffusion’s vanilla text-to-image and image-to-image pipelines to generate synthetic images. Next, we leverage Textual Inversion to learn the lesion embeddings and then generate synthetic images from these embeddings, with text-to-image and image-to-image pipelines. Since image-to-image outperforms text-to-image in both vanilla SD and Textual Inversion generation, we focus on image-to-image generation after fine-tuning the diffusion model using LoRA, to investigate if optimizing the diffusion model can benefit the generation even more. We train a VGG-16 using only the synthetic images and then compare these generation strategies with ours, as shown in Table 5.

Overall, training classifiers on synthetic images generated by text-to-image models proves less effective than employing image-to-image techniques, underscoring the importance of visual cues in augmenting skin lesion classification. Additionally, using off-the-shelf models for image generation yields less improvement than training strategies such as Textual Inversion and LoRA, regardless of whether the target is a minority or majority group. Finally, the combination of Textual Inversion and LoRA results in the highest accuracy, thereby validating the practicality of our design which integrates these two strategies. This improvement can be explained by the model’s enhanced ability to associate the fine visual cues of the lesion with the learned textual tokens.

Since a direct comparison with existing related works is challenging due to the uncertain use of data, this ablation study can serve as an indirect comparison. As introduced previously, related works leveraged off-the-shelf diffusion models such as DALL·E or fine-tuned a Stable Diffusion model for text-to-image generation. Our results demonstrate that utilizing the dual guidance of visual cues and text



**Fig. 2.** Examples of synthetic images generated by a model trained exclusively on light-skinned images, using prompts describing dark skin types.

test	txt2img		img2img			
	vanilla	ti	vanilla	ti	lora	ti+lora
light (56)	18.80	35.36	48.21	46.43	52.00	53.57
dark (56)	21.22	44.64	69.64	71.43	73.21	79.57

**Table 5.** Classification accuracy with various generation strategies for the first setting. Here, “vanilla” stands for Stable Diffusion’s original text-to-image and image-to-image pipelines, “ti” for Textual Inversion, and “lora” for LoRA.

prompts via fine-tuning diffusion models can maximize the potential of diffusion-based augmentation and enhance the diagnosis for minority groups.

## 5 Conclusion

In this work, we present an effective diffusion-based augmentation framework that consistently improves classification results for the minority group, even when training the classifier exclusively with synthetic images. This improvement is observed regardless of the availability of reference data from the minority group in the training set. The ablation study also validates that our framework’s dual-guidance generation approach successfully learns novel lesion concepts previously unfamiliar to the diffusion models. A practical takeaway from this study is that, even in cases of data scarcity, existing data and diffusion models can still provide valuable insights, maximizing information usage and achieving better performance. In the future, we plan to apply this technique to other medical datasets characterized by significant differences in group sizes. Additionally, as we used all synthetic images generated for each setting without any filtering mechanism, we also aim to investigate which types of synthetic data are useful for lesion diagnosis and how to generate them.

**Acknowledgments.** This work was partly supported by the NSF EPSCoR-Louisiana Materials Design Alliance (LAMDA) program #OIA-1946231 and partly by the Harold L. and Heather E. Jurist Center of Excellence for Artificial Intelligence at Tulane University.



## References

1. Akrouf, M., Gyepesi, B., Holló, P., Poór, A., Kincsó, B., Solis, S., Cirone, K., Kawahara, J., Slade, D., Abid, L., Kovács, M., Fazekas, I.: Diffusion-based data augmentation for skin disease classification: Impact across original medical datasets to fully synthetic images (2023)
2. Brinker, T., Hekler, A., Enk, A., Klode, J., Hauschild, A., Berking, C., Schilling, B., Haferkamp, S., Utikal, J., Kalle, C., Fröhling, S., Weichenthal, M.: A convolutional neural network trained with dermoscopic images performed on par with 145 dermatologists in a clinical melanoma image classification task. *European Journal of Cancer* **111**, 148–154 (03 2019). <https://doi.org/10.1016/j.ejca.2019.02.005>
3. Coustasse, A., Sarkar, R., Abodunde, B., Metzger, B.J., Slater, C.M.: Use of tele-dermatology to improve dermatological access in rural areas. *Telemedicine journal and e-health : the official journal of the American Telemedicine Association* (2019)
4. Daneshjou, R., Vodrahalli, K., Liang, W., Novoa, R.A., Jenkins, M., Rotemberg, V., Ko, J., Swetter, S.M., Bailey, E.E., Gevaert, O., Mukherjee, P., Phung, M., Yekrang, K., Fong, B., Sahasrabudhe, R., Zou, J., Chiou, A.: Disparities in dermatology ai performance on a diverse, curated clinical image set. *Science Advances* (2022)
5. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
6. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J.M., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**, 115–118 (2017)
7. Fitzpatrick, T.B.: The validity and practicality of sun-reactive skin types i through vi. *Archives of dermatology* **124**(6), 869–871 (1988)
8. Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=NAQvF08TcyG>
9. Ghorbani, A., Natarajan, V., Coz, D., Liu, Y.: DermGAN: Synthetic Generation of Clinical Skin Images with Pathology. In: Dalca, A.V., McDermott, M.B., Alsentzer, E., Finlayson, S.G., Oberst, M., Falck, F., Beaulieu-Jones, B. (eds.) *Proceedings of the Machine Learning for Health NeurIPS Workshop. Proceedings of Machine Learning Research*, vol. 116, pp. 155–170. PMLR (13 Dec 2020), <https://proceedings.mlr.press/v116/ghorbani20a.html>
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)
11. Groh, M., Harris, C., Soenksen, L., Lau, F., Han, R., Kim, A., Koochek, A., Badri, O.: Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1820–1828 (2021)
12. 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)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
14. Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: *International Con-*

- ference on Learning Representations (2022), <https://openreview.net/forum?id=nZeVKeeFYf9>
15. Ktena, I., Wiles, O., Albuquerque, I., Rebuffi, S.A., Tanno, R., Roy, A.G., Azizi, S., Belgrave, D., Kohli, P., Karthikesalingam, A., Cemgil, T., Goyal, S.: Generative models improve fairness of medical classifiers under distribution shifts (2023)
  16. Liu, Y., Jain, A., Eng, C., Way, D.H., Lee, K., Bui, P., Kanada, K., de Oliveira Marinho, G., Gallegos, J., Gabriele, S., Gupta, V., Singh, N., Natarajan, V., Hofmann-Wellenhof, R., Corrado, G.S., Peng, L.H., Webster, D.R., Ai, D., Huang, S.J., Liu, Y., Dunn, R.C., Coz, D.: A deep learning system for differential diagnosis of skin diseases. *Nature Medicine* **26**(6), 900–908 (Jun 2020)
  17. von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Nair, D., Paul, S., Berman, W., Xu, Y., Liu, S., Wolf, T.: Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers> (2022)
  18. Qin, Z., Liu, Z., Zhu, P., Xue, Y.: A gan-based image synthesis method for skin lesion classification. *Computer Methods and Programs in Biomedicine* **195**, 105568 (2020)
  19. Rezk, E., Eltorki, M., El-Dakhkhni, W., et al.: Improving skin color diversity in cancer detection: deep learning approach. *JMIR Dermatology* **5**(3), e39143 (2022)
  20. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10684–10695 (June 2022)
  21. 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 (2023)
  22. Sagers, L.W., Diao, J.A., Groh, M., Rajpurkar, P., Adamson, A., Manrai, A.K.: Improving dermatology classifiers across populations using images generated by large diffusion models. In: *NeurIPS 2022 Workshop on Synthetic Data for Empowering ML Research* (2022), <https://openreview.net/forum?id=Vzdbjtz6Tys>
  23. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
  24. Wang, J., Zhang, Y., Ding, Z., Hamm, J.: Achieving reliable and fair skin lesion diagnosis via unsupervised domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5157–5166 (2024)
  25. Wu, B., Xu, C., Dai, X., Wan, A., Zhang, P., Yan, Z., Tomizuka, M., Gonzalez, J., Keutzer, K., Vajda, P.: Visual transformers: Token-based image representation and processing for computer vision (2020)