

Controllable Skin Synthesis via Lesion-Focused Vector Autoregression Model

Jiajun Sun^{1,2}, Zhen Yu^{3,4}, Siyuan Yan^{3,5}, Jason J. Ong^{1,2,6}, Zongyuan Ge^{3,4,*},
and Lei Zhang^{1,2,7,8,*}

¹ School of Translational Medicine, Faculty of Medicine, Nursing and Health Sciences, Monash University, Melbourne, Australia

² Melbourne Sexual Health Centre, Alfred Health, Melbourne, Australia

³ AIM for Health Lab, Monash University, Melbourne, VIC, Australia

⁴ Faculty of IT, Monash University, Melbourne, VIC, Australia

⁵ Faculty of Engineering, Monash University, Melbourne, VIC, Australia

⁶ Faculty of Infectious and Tropical Diseases, London School of Hygiene and Tropical Medicine, London, United Kingdom

⁷ China-Australia Joint Research Center for Infectious Diseases, School of Public Health, Xi'an Jiaotong University Health Science Center, Xi'an, Shaanxi, China

⁸ Phase I clinical trial research ward, The Second Affiliated Hospital of Xi'an Jiaotong University, Xi'an, Shaanxi, China.

{Jiajun.Sun, Zhen.Yu, Siyuan.Yan, Jason.Ong, Zongyuan.Ge,
Lei.Zhang1}@monash.edu

Abstract. Skin images from real-world clinical practice are often limited, resulting in a shortage of training data for deep-learning models. While many studies have explored skin image synthesis, existing methods often generate low-quality images and lack control over the lesion's location and type. To address these limitations, we present **LF-VAR**, a model leveraging quantified lesion measurement scores and lesion type labels to guide the clinically relevant and controllable synthesis of skin images. It enables controlled skin synthesis with specific lesion characteristics based on language prompts. We train a multiscale lesion-focused Vector Quantised Variational Auto-Encoder (VQVAE) to encode images into discrete latent representations for structured tokenization. Then, a Visual AutoRegressive (VAR) Transformer trained on tokenized representations facilitates image synthesis. Lesion measurement from the lesion region and types as conditional embeddings are integrated to enhance synthesis fidelity. Our method achieves the best overall FID score (average 0.74) among seven lesion types, improving upon the previous state-of-the-art (SOTA) by **6.3%**. The study highlights our controllable skin synthesis model's effectiveness in generating high-fidelity, clinically relevant synthetic skin images. Our framework code is available at <https://github.com/echosun1996/LF-VAR>.

Keywords: Skin Image Synthesis · Vector Autoregression · Dermatology.

* Corresponding author.

1 Introduction

High-quality skin images are essential for training and validating medical artificial intelligence (AI) models. However, skin images from real-world clinical practice often suffer from scarcity and long-tail imbalances, leading to biased outcomes or algorithmic unfairness [21]. Furthermore, privacy concerns, legal restrictions, and financial burdens pose substantial challenges in acquiring diverse, high-quality skin images [8, 7].

Generative models have emerged as a potential solution to address these data-related challenges. For instance, SkinDiff [11] employs a latent diffusion framework to generate labelled skin lesion images, while Derm-T2IM [6] leverages diffusion models with lesion-specific textual prompts to generate dermoscopic images. However, diffusion models are limited by their dependence on large-scale training data and computationally intensive synthesis processes [19, 3], making them less feasible for resource-constrained clinical settings.

Recent advances in autoregressive models, particularly Vector AutoRegressive (VAR) Transformers, have demonstrated competitive performance in general image synthesis tasks while requiring fewer computational resources and less training data than diffusion models [22]. However, our evaluation of the VAR model in skin image synthesis revealed two critical limitations: First, existing implementations lack mechanisms to guide synthesis using diagnostically relevant features, resulting in generated images with limited pathological diversity and fidelity. Second, without explicit separation of lesions and normal skin regions, VAR-generated images exhibit unnatural transitions and background noise, reducing their realism.

To address these challenges, we propose *LF-VAR*, a controllable synthesis model that leverages masks and lesion type as prompts to enable precise and flexible skin lesion synthesis. Additionally, *LF-VAR* integrates quantified lesion measurements—such as shape and texture—to guide our model over clinically relevant features, ensuring high-fidelity synthetic images with well-preserved pathological characteristics. The primary contributions of our work are:

1. We propose a lightweight encoder that transforms quantified lesion measurements into discriminative latent representations, which enables precise control over clinically relevant features and allows high-fidelity synthesis.
2. We implement a unified model that incorporates segmentation masks and lesion type, enabling controlled synthesis of lesions across different lesion categories. This design allows for the flexible generation of diverse pathological presentations while maintaining consistency with clinical descriptors.
3. We integrate pixel-level segmentation into the generative pipeline to focus lesion synthesis while minimizing background region artifacts. It enables high-quality image synthesis by maintaining a clear distinction between lesions and normal skin regions.
4. We conduct extensive validation indicating robust synthesis across intra-class, inter-class, and cross-dataset scenarios. Our evaluation confirms the model’s effect in high-fidelity clinically relevant skin synthetic.




Fig. 1. The framework of *LF-VAR*. A structured VQ-VAE encoded images into K token maps, shared with the VAR, and was trained with lesion-focused loss. Lesion measurement scores were extracted from lesion areas and encoded as $\mathbf{F}_q$, combined with lesion type embedding $\mathbf{S}$, and incorporated into the VAR as a conditional token to enable controlled intra-class skin lesion synthesis. Inter-class synthesis was achieved through the VAR model by querying the lesion type average measurements codebook.

2 Method

The proposed synthesis model, as shown in Fig.1 consisted of two distinct stages: (1) a multiscale lesion-focused Vector Quantised Variational Auto-Encoder (VQ-VAE) [16], which was employed to acquire compact discrete image representations, and (2) a VAR [22] model trained on tokenised representations to facilitate image synthesis. In the second stage, measurement scores extracted from the skin lesion region were incorporated into the transformer, along with lesion types as conditional embeddings, to improve the fidelity of the generated images.

2.1 Multi-Scale Lesion-Focused VQ-VAE

A multi-scale lesion-focused VQ-VAE was trained to encode images into discrete latent representations, enabling structured and hierarchical tokenization. Given an input image $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$, the encoder $E(\cdot)$ extracted feature maps, yielding a latent feature map $f = E(\mathbf{I})$, where $f \in \mathbb{R}^{h \times w \times d}$. A multi-scale quantization process was subsequently applied to discretize the latent representation:

$$\mathbf{R} = Q(f) = \{R_k\}_{k=1}^K, \quad R_k \in \mathbb{Z}^{h_k \times w_k}, \quad (1)$$

where $Q(\cdot)$ denoted the quantization function, K represented the number of scales, and each R_k was a discrete token grid at scale k . The decoder networks $D(\cdot)$ and $\hat{D}(\cdot)$ were responsible for reconstructing the feature map and the final

image, respectively, where the reconstructed feature map was given by $\hat{f} = D(\mathbf{R})$ and the final reconstructed image by $\hat{\mathbf{I}} = \hat{D}(\hat{f})$.

To enhance reconstruction fidelity, a loss function was introduced to enforce both pixel-level and perceptual constraints while ensuring feature consistency in non-masked regions across different scales. Overall loss function was defined as:

$$\mathcal{L} = \|\mathbf{I} - \hat{\mathbf{I}}\|_2 + \sum_{k=1}^K \|(1 - M_k) \odot (R_k - \hat{R}_k)\|_2 + \|f - \hat{f}\|_2 + \lambda_P \mathcal{L}_P(\hat{\mathbf{I}}) + \lambda_G \mathcal{L}_G(\hat{\mathbf{I}}), \quad (2)$$

where $M_k \in \{0, 1\}^{h_k \times w_k}$ was the lesion mask at scale k . The second term was lesion-focused loss, enforcing similarity between real and reconstructed feature tokens in non-masked areas across all scales. $\mathcal{L}_P$ represented a perceptual loss, while $\mathcal{L}_G$ represented a discriminative loss. λ_P and λ_G were the loss weights.

2.2 Lesion Measurement Embedded VAR Transformer

In clinical dermatology, quantified lesion measurements—structured scores that captured morphological characteristics such as lesion texture and shape—were essential for both diagnosis and prognosis [21, 14, 13]. These structured scores, derived from mask regions rather than full images, offered robust feature representations. Despite their clinical importance, the potential of these structured measurements to guide medical image synthesis remained largely unexplored.

To bridge this gap, we incorporated lesion measurement scores as conditioning signals for controlled skin lesion synthesis. Specifically, during intra-class synthesis, lesion measurement scores ε_r were extracted from the original images using PyRadiomics [9] formulated as $\varepsilon_r = E_{\text{ext}}(\mathbf{I}, \mathbf{M})$. These measurement scores included shape, histogram, texture, and clinical attributes.

To enable autoregressive modelling of discrete tokens, a VAR Transformer was employed, conditioning synthesis on learned embeddings. Along with extracted multi-scale tokens, measurement scores were incorporated via an encoding function, $\mathbf{F}_q = E_q(\varepsilon_r)$. The measurement scores ε_r were transformed via a linear projection, followed by layer normalization and a SiLU activation function.

For inter-class synthesis, an additional codebook of class-average measurements was maintained. During the intra-class synthesis, ε_r was used to update the corresponding skin lesion entry in the codebook by computing the average measurements. During the inter-class synthesis, the query feature $\hat{\mathbf{F}}_q$ was obtained by referencing this codebook.

A class embedding $\mathbf{S}$ was introduced to incorporate class-specific information for the generative process. The combined token sequence was defined as $\mathbf{T}_0 = [\mathbf{S}, \mathbf{F}_r]$. The VAR Transformer autoregressively predicted the subsequent tokens, denoted as $\hat{\mathbf{R}}$, based on the input sequence $\mathbf{T}_0$ and the reference sequence $\mathbf{R}$. A “next-token prediction” generated the higher-scale skin lesion representations based on the combined tokens and the corresponding lower-scale tokens. The training objective was a standard cross-entropy loss over the token space.




Fig. 2. Comparison of synthesis samples by baseline and our method, with two samples from different lesions of each model.

3 Experiments and Results

3.1 Experiments Setup

Dataset. In this study, the HAM10000 [23] dataset was utilized to train the *LF-VAR* model. The dataset consisted of dermatoscopy images annotated with seven skin lesion categories: actinic keratoses (AKIEC, 327 images), basal cell carcinoma (BCC, 514), benign keratosis (BKL, 1,099), dermatofibroma (DF, 118), melanoma (MEL, 1,113), melanocytic nevi (NV, 6,705), and vascular lesions (VASC, 142). All images were resized to a resolution of 512×512 pixels and split into training and testing sets with a 4:1 ratio. For cross-dataset performance evaluation, images from the ISIC2017 [4] and Dermofit [2] datasets that matched the lesion categories in the HAM10000 dataset were selected.

Baselines. Several generative methods were selected as baselines, including MAGE [12], MAGE Adapter [25], Derm T2IM [6], VAR [22], and Diffusion [18]. MAGE and MAGE Adapter were unconditional image generation models. Derm T2IM and VAR utilized text as prompts to guide image synthesis. The Diffusion and VAR model possessed inpainting capabilities, allowing it to generate content within a specified masked region while completing the missing area.

Task and metric. The proposed method was evaluated on three tasks: (1) Intra-class synthesis, to assess the model’s ability to replicate patterns from the training set; (2) Inter-class synthesis, to evaluate its generalizability across different skin lesion types; (3) Cross-dataset evaluation, to test the model’s robustness and generalization on external datasets, and (4) Validation on the downstream classification tasks.

Two metrics were employed: (1) Fréchet Inception Distance (FID) [10], which measured the distance between the feature distributions of real and synthesised images, with a lower FID indicating higher fidelity; (2) Inception Score (IS) [20], which evaluated the quality and diversity of the generated images, with a higher




Fig. 3. Feature analysis and model comparison. (a) Feature distribution across seven skin lesion types. (b) Feature distribution comparison between real and synthetic in melanoma. (c) Synthesis images feature distance and FID comparison among models.

IS reflecting improved variation. In this study, we primarily focus on FID as it provides a more reliable measure of fidelity by comparing feature distributions between real and synthesized images.

Implementation details. The models were trained using the AdamW optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.95$), a weight decay of 0.05, a learning rate of 1×10^{-3} , and a batch size of 35 for 200 epochs. All experiments ran on an NVIDIA A100 GPU (40 GB) in a high-performance computing environment. We employed a multi-scale quantization autoencoder[22, 1] and GPT-2 decoder-only transformer [5, 17] which keeps consistent with VAR.

3.2 Results

Intra-class synthesis. Fig.2 showed randomly selected skin images generated by different methods for visual comparison. MAGE [12] and MAGE Adapter [25] often produced blurred or distorted images with lower visual quality. Derm T2IM [6] and Diffusion [18] models generated higher-quality results, though inconsistencies were noticeable, especially for melanoma. VAR [22] produced visually realistic images but failed to capture key lesion characteristics, with inconsistent lesion colour, shape, and unnatural transitions. In contrast, our method improved skin synthesis quality through lesion measurement scores and reduced artifacts with a well-designed multi-scale lesion-focused VQ-VAE, resulting in more realistic skin lesion images.

Fig.3 presented a feature space analysis of real and generated images. The t-distributed Stochastic Neighbor Embedding (t-SNE) [15] visualization showed distinct clusters for each skin lesion type. Feature distance comparisons and FID scores across lesion types indicated that our method achieved better feature consistency and synthesis quality compared to baseline models. Table 1 presented the quantitative evaluation of synthesis performance using FID and IS metrics. Our method achieved the best overall FID score (average 0.74), improving upon the previous state-of-the-art (SOTA) by 6.3%. Derm T2IM showed high IS across most categories (4.28), due to the introduction of artifacts into normal skin regions, which increased diversity and inflated the IS metric. In contrast, our

Table 1. Performance comparison on the test dataset. **Bold** indicates the best FID, while underlined indicates the second-best.

Method	Prompt	Matrix	AKIEC	BCC	BKL	DF	MEL	NV	VASC	Average
MAGE [12]	$\times$	IS $\uparrow$	3.08	3.18	3.01	2.96	2.91	2.98	3.08	3.03
			$\pm 0.13 \pm 0.12 \pm 0.15 \pm 0.21 \pm 0.16 \pm 0.14 \pm 0.11$							± 0.15
		FID $\downarrow$	12.92	11.86	9.84	12.18	8.49	14.97	16.34	12.37
MAGE Adapter [25]	$\times$	IS $\uparrow$	2.49	1.79	1.98	2.84	2.38	2.66	2.67	2.40
			$\pm 0.05 \pm 0.09 \pm 0.05 \pm 0.14 \pm 0.09 \pm 0.10 \pm 0.16$							± 0.10
		FID $\downarrow$	2.33	7.99	2.65	3.09	5.82	3.24	3.28	4.06
Derm T2IM [6]	Text	IS $\uparrow$	3.74	4.23	3.94	5.03	3.59	4.74	4.69	4.28
			$\pm 0.19 \pm 0.17 \pm 0.39 \pm 0.31 \pm 0.29 \pm 0.30 \pm 0.46$							± 0.30
		FID $\downarrow$	6.52	6.02	3.32	5.31	4.74	6.26	5.69	5.41
Diffusion [18]	Mask	IS $\uparrow$	3.90	4.88	3.12	4.04	4.22	3.96	3.10	3.89
			$\pm 0.20 \pm 0.25 \pm 0.12 \pm 0.18 \pm 0.30 \pm 0.25 \pm 0.09$							± 0.20
		FID $\downarrow$	1.22	<u>1.47</u>	0.78	1.15	0.90	1.57	0.64	1.11
VAR [22]	Text-Mask	IS $\uparrow$	3.09	2.57	3.27	2.58	2.76	2.37	2.48	2.73
			$\pm 0.14 \pm 0.12 \pm 0.19 \pm 0.08 \pm 0.11 \pm 0.10 \pm 0.15$							± 0.13
	Mask	FID $\downarrow$	<u>0.74</u>	1.58	0.44	<u>0.76</u>	0.44	0.56	<u>1.04</u>	<u>0.79</u>
Ours	Text-Mask	IS $\uparrow$	3.27	2.41	3.26	2.34	2.84	2.63	3.92	2.95
			$\pm 0.13 \pm 0.09 \pm 0.12 \pm 0.11 \pm 0.08 \pm 0.13 \pm 0.08$							± 0.10
	Mask	FID $\downarrow$	0.60	1.04	<u>0.69</u>	0.63	<u>0.45</u>	<u>0.67</u>	1.09	0.74

method maintained an IS of 2.95 while preserving lesion fidelity and minimizing artifacts, as reflected in its lowest FID scores.

Inter-class synthesis. We evaluated the inter-class synthesis performance of our model using different text-based lesion categories as prompts. The results were presented in Fig. 4. As shown in the confusion matrix, all lesion categories achieved the best synthesis performance when generating VASC, with an average FID of 0.90 (standard deviation [std], 0.13), followed by DF (mean, 0.96; std, 0.27). The comparative figure further showed that our model preserved a clear background and smooth lesion boundaries while maintaining consistent lesion characteristics across columns, ensuring the capture of distinct lesion features in inter-class synthesis.

Cross-dataset synthesis. Table 2 presented the evaluation results on two external datasets. Our method achieved the best average performance across both datasets. Specifically, on the ISIC2017 dataset, it obtained the lowest FID (1.70), outperforming all other methods across most categories, particularly in BKL (1.94) and MEL (1.99). On the Dermofit dataset, our method achieved the best average FID (1.71), with the lowest FID in AKIEC (1.29), BCC (0.56), and VASC (0.38). These results highlighted the synthesis capability of our model compared to existing approaches, showing consistent performance across diverse datasets.

Downstream task validation. Following the established protocol by Philipp et al. [24], our downstream classification uses a ResNet-50 classifier for 7 lesion types. The baseline achieved a mean recall of 0.692. Adding a weighted random sampler improved it to 0.715. When using augmented training data with our synthetic images (balanced to 500 images per class), the recall increased to 0.771 (+0.079, +11.4% from baseline; +0.056, +7.8% from weighted sampling).




Fig. 4. Inter-class synthesis comparison and FID matrix. (a) Inter-class synthesis samples for seven lesion types. The first two columns show the original images and masks, followed by synthesis images. (b) FID confusion matrix for inter-class synthesis.

Table 2. FID performance comparison on cross-dataset synthesis.

Dataset	Method	AKIEC	BCC	BKL	DF	MEL	NV	VASC	Average
ISIC2017 [4]	MAGE [12]	-	-	8.04	-	6.99	8.34	-	7.94
	MAGE Adapter [25]	-	-	5.89	-	8.19	9.75	-	7.79
	Derm T2IM [6]	-	-	4.22	-	6.01	<u>3.88</u>	-	4.71
	VAR [22]	-	-	2.49	-	<u>2.13</u>	4.21	-	<u>2.94</u>
	Diffusion [18]	-	-	<u>2.36</u>	-	7.20	12.98	-	7.51
	Ours	-	-	1.94	-	1.99	1.18	-	1.70
Dermofit [2]	MAGE [12]	11.01	11.77	10.42	7.80	11.32	9.00	7.91	9.89
	MAGE Adapter [25]	6.88	16.02	8.76	5.31	8.24	15.90	6.89	9.71
	Derm T2IM [6]	8.29	10.69	9.08	9.20	13.48	6.78	7.27	9.26
	VAR [22]	4.94	5.40	4.98	5.78	4.64	8.25	9.43	6.20
	Diffusion [18]	<u>2.20</u>	<u>3.88</u>	2.23	2.67	12.73	2.38	<u>1.90</u>	<u>4.00</u>
	Ours	1.29	0.56	2.23	<u>3.92</u>	0.39	<u>3.23</u>	0.38	1.71

This demonstrates the practical utility of our synthetic data for improving classification performance.

Ablation study. To validate the effectiveness of lesion focus VQ-VAE and measurement embedding in enhancing VAR, we systematically integrated specific modules into the baseline. We sequentially introduced the following components: **LF**: a multi-scale lesion focus VQ-VAE; **FM**: a fixed measurement embedding where a constant measurement value was applied across all lesions; and **AM**: an adaptive measurement embedding that maintained a class-specific average measurement codebook and dynamically assigned measurements based on lesion type.

Table 3 summarised the experimental study results, demonstrating a progressive improvement in synthesis performance with the incorporation of different modules. Integrating LF into the baseline led to an average performance gain of 3.7%, while combining LF and AM led to the best overall performance, achieving the lowest average FID (0.74) and the highest IS (2.95 ± 0.10). Compared to

Table 3. Ablation study results. **Bold** indicates the best FID, while underlined represents the second-best.

Setting	Metric	AKIEC	BCC	BKL	DF	MEL	NV	VASC	Average
Baseline	IS $\uparrow$	3.09	2.57	3.27	2.58	2.76	2.37	2.48	2.73
		± 0.14	± 0.12	± 0.19	± 0.08	± 0.11	± 0.10	± 0.15	± 0.13
	FID $\downarrow$	0.74	1.58	0.44	0.76	0.44	0.56	1.04	0.79
Baseline + LF	IS $\uparrow$	3.42	2.52	3.47	3.06	3.01	2.67	2.35	2.93
		± 0.10	± 0.10	± 0.26	± 0.11	± 0.12	± 0.17	± 0.08	± 0.13
	FID $\downarrow$	<u>0.70</u>	<u>1.23</u>	<u>0.47</u>	0.62	0.54	0.74	<u>1.00</u>	<u>0.76</u>
Baseline + LF + FM	IS $\uparrow$	2.69	2.05	2.95	2.65	2.42	2.25	1.77	2.40
		± 0.08	± 0.07	± 0.17	± 0.11	± 0.09	± 0.12	± 0.03	± 0.10
	FID $\downarrow$	1.14	2.09	0.73	0.72	1.16	<u>0.62</u>	0.91	1.05
Baseline + LF + AM	IS $\uparrow$	3.27	2.41	3.26	2.34	2.84	2.63	3.92	2.95
		± 0.13	± 0.09	± 0.12	± 0.11	± 0.08	± 0.13	± 0.08	± 0.10
	FID $\downarrow$	0.60	1.04	0.69	<u>0.63</u>	<u>0.45</u>	0.67	1.09	0.74

the baseline, combining LF and AM improved performance by 6.3%, surpassing the baseline with LF by 2.6%.

4 Conclusion

In this paper, we introduced *LF-VAR*, a controllable skin lesion synthesis model that integrates lesion measurement scores with a lesion-focused VQ-VAE and VAR model. By leveraging structured lesion measurements as conditional signals, our approach enables precise control over key lesion characteristics while maintaining high synthesis fidelity. The results demonstrate the effectiveness of integrating clinically relevant features into generative models, highlighting the potential of *LF-VAR* for enhancing skin image synthesis in medical applications.

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

References

1. Bai, Y., Geng, X., Mangalam, K., Bar, A., Yuille, A., Darrell, T., Malik, J., Efros, A.A.: Sequential modeling enables scalable learning for large vision models (2023), <https://arxiv.org/abs/2312.00785>
2. Ballerini, L., Fisher, R.B., Aldridge, B., Rees, J.: A color and texture based hierarchical k-nn approach to the classification of non-melanoma skin lesions. *Color medical image analysis* pp. 63–86 (2013)
3. Chen, M., Mei, S., Fan, J., Wang, M.: Opportunities and challenges of diffusion models for generative ai. *National Science Review* **11**(12), nwae348 (10 2024). <https://doi.org/10.1093/nsr/nwae348>
4. Codella, N.C.F., Gutman, D., Celebi, M.E., Helba, B., Marchetti, M.A., Dusza, S.W., Kallou, A., Liopyris, K., Mishra, N., Kittler, H., Halpern, A.: Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (ISBI), hosted by the international skin imaging collaboration (ISIC). <https://doi.org/10.48550/arXiv.1710.05006>

5. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis (2021), <https://arxiv.org/abs/2012.09841>
6. Farooq, M.A., Yao, W., Schukat, M., Little, M.A., Corcoran, P.: Derm-t2im: Harnessing synthetic skin lesion data via stable diffusion models for enhanced skin disease classification using vit and cnn. In: 2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). p. 1–5. IEEE (Jul 2024). <https://doi.org/10.1109/embc53108.2024.10781852>
7. Finnane, A., Curiel-Lewandrowski, C., Wimberley, G., Caffery, L., Katragadda, C., Halpern, A., Marghoob, A.A., Malvey, J., Kittler, H., Hofmann-Wellenhof, R., Abraham, I., Soyer, H.P., behalf of the International Society of Digital Imaging of the Skin (ISDIS) for the International Skin Imaging Collaboration (ISIC), O.: Proposed technical guidelines for the acquisition of clinical images of skin-related conditions. *JAMA Dermatology* **153**(5), 453–457 (05 2017). <https://doi.org/10.1001/jamadermatol.2016.6214>
8. Gordon, E.R., Trager, M.H., Kontos, D., Weng, C., Geskin, L.J., Dugdale, L.S., Samie, F.H.: Ethical considerations for artificial intelligence in dermatology: a scoping review. *British Journal of Dermatology* **190**(6), 789–797 (02 2024). <https://doi.org/10.1093/bjd/ljae040>
9. van Griethuysen, J.J., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R.G., Fillion-Robin, J.C., Pieper, S., Aerts, H.J.: Computational radiomics system to decode the radiographic phenotype. *Cancer Research* **77**(21), e104–e107 (10 2017). <https://doi.org/10.1158/0008-5472.CAN-17-0339>
10. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium (2018)
11. Jing, X., Yang, S., Zhou, H., Wang, G., Mao, K.: Skindiff: A novel data synthesis method based on latent diffusion model for skin lesion segmentation. In: Advanced Intelligent Computing Technology and Applications: 20th International Conference, ICIC 2024, Tianjin, China, August 5–8, 2024, Proceedings, Part VIII. p. 179–191. Springer-Verlag, Berlin, Heidelberg (2024). https://doi.org/10.1007/978-981-97-5603-2_15
12. Li, T., Chang, H., Mishra, S., Zhang, H., Katabi, D., Krishnan, D.: Mage: Masked generative encoder to unify representation learning and image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2142–2152 (2023)
13. Li, Y., Ammari, S., Lawrance, L., Quillent, A., Assi, T., Lassau, N., Chouzenoux, E.: Radiomics-based method for predicting the glioma subtype as defined by tumor grade, idh mutation, and 1p/19q codeletion. *Cancers* **14**(7), 1778 (2022)
14. Liu, Z., Wang, S., Dong, D., Wei, J., Fang, C., Zhou, X., Sun, K., Li, L., Li, B., Wang, M., et al.: The applications of radiomics in precision diagnosis and treatment of oncology: opportunities and challenges. *Theranostics* **9**(5), 1303 (2019)
15. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(11) (2008)
16. van den Oord, A., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning (2018)
17. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019)
18. 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)

19. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2022)
20. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., Chen, X.: Improved techniques for training gans. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016)
21. Staal, F.C., van der Reijnd, D.J., Taghavi, M., Lambregts, D.M., Beets-Tan, R.G., Maas, M.: Radiomics for the prediction of treatment outcome and survival in patients with colorectal cancer: A systematic review. *Clinical Colorectal Cancer* **20**(1), 52–71 (2021). <https://doi.org/10.1016/j.clcc.2020.11.001>
22. Tian, K., Jiang, Y., Yuan, Z., PENG, B., Wang, L.: Visual autoregressive modeling: Scalable image generation via next-scale prediction. In: Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C. (eds.) *Advances in Neural Information Processing Systems*. vol. 37, pp. 84839–84865. Curran Associates, Inc. (2024)
23. Tschandl, P.: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions (2018). <https://doi.org/10.7910/DVN/DBW86T>
24. Tschandl, P., Rinner, C., Apalla, Z., Argenziano, G., Codella, N., Halpern, A., Janda, M., Lallas, A., Longo, C., Malvehy, J., et al.: Human–computer collaboration for skin cancer recognition. *Nature medicine* **26**(8), 1229–1234 (2020)
25. Yuan, Z., Fang, Z., Huang, Z., Wu, F., Yao, Y.F., Li, Y.: Adapting Pre-trained Generative Model to Medical Image for Data Augmentation . In: *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*. vol. LNCS 15005. Springer Nature Switzerland (October 2024)