

Automated Disentangling Analysis of Skin Colour for Lesion Images

Wenbo Yang^{*[0009-0000-7528-1925]}, Eman Rezk^[0000-0002-2531-0799], Walaa M. Moursi^[0000-0002-0113-9309], and Zhou Wang^[0000-0003-4413-4441]
{w243yang,e2rezk,walaa.moursi,z70wang}@uwaterloo.ca

University of Waterloo, Waterloo, ON, Canada, N2L 3G1

Abstract. Machine-learning models applied to skin images often have degraded performance when the skin colour captured in images (SCCI) differs between training and deployment. These discrepancies arise from a combination of entangled environmental factors (e.g., illumination, camera settings) and intrinsic factors (e.g., skin tone) that cannot be accurately described by a single “skin tone” scalar – a simplification commonly adopted by prior work. To mitigate such colour mismatches, we propose a skin-colour disentangling framework that adapts disentanglement-by-compression to learn a structured, manipulable latent space for SCCI from unlabelled dermatology images. To prevent information leakage that hinders proper learning of dark colour features, we introduce a randomized, mostly monotonic decolourization mapping. To suppress unintended colour shifts of localized patterns (e.g., ink marks, scars) during colour manipulation, we further propose a geometry-aligned post-processing step. Together, these components enable faithful counterfactual editing and answering an essential question: “What would this skin condition look like under a different SCCI?”, as well as direct colour transfer between images and controlled traversal along physically meaningful directions (e.g., blood perfusion, camera white balance), enabling educational visualization of skin conditions under varying SCCI. We demonstrate that dataset-level augmentation and colour normalization based on our framework achieve competitive lesion classification performance. Ultimately, our work promotes equitable diagnosis through creating diverse training datasets that include different skin tones and image-capturing conditions.

Keywords: Skin Colour Disentanglement · Counterfactual Image Synthesis · Data Augmentation

1 Introduction

Computer-aided dermatology diagnosis systems have improved significantly in recent years. However, similar to human experts, their performance degrades when training and deployment data exhibit mismatched skin colours captured

^{*} Corresponding author




Fig. 1. Our model is capable of changing the captured skin colour after training.




Fig. 2. Trajectories with physical meanings can be found in the latent space.

in images (SCCI). Benmalek et al. [3] reports accuracy gaps reaching 10%–30% between light and dark skin presentations. The mismatch of SCCI is caused by multiple entangled factors, including intrinsic subject properties (e.g., skin tone), and environmental factors (e.g., illumination, camera white balance, and other device settings) [10, 24].

Existing mitigation strategies for colour mismatch fall into two categories. First, most recent works in dermatology focus on augmenting the training set using augmented images [1] or partially or fully synthetic images [17, 8]. Despite not explicitly modelling other race-specific lesion characteristics, these intrinsic-factor-centric methods show promising performance improvement, indicating that biases removed by skin-tone transfer significantly outweigh those introduced by these methods. However, these model usually consider skin colour on individual images on an *ad hoc* basis without a structured colour model, of-




Fig. 3. Architecture of the skin colour disentangling model

fering limited control of the result and distribution matching of datasets. The single-value scalar or categorical skin-tone descriptors (such as FST [7] and MST [16]) used by these methods have been shown to have very poor correlation with the skin colour shown in photos [10, 24], because it does not reflect other factors (including blood perfusion and environmental factors). Second, colour normalization aims at mitigating the mismatch by converting each image to a “standard” colour. While such a method has been widely used in histology [4, 25, 23], only limited works have been proposed for skin images [13, 21, 20], possibly due to the more complex colour formation process. Currently, these methods focus on environmental factors without explicitly considering the intrinsic factors that affect the SCCI. Moreover, educators and future physicians need a tool for visualizing the same skin condition under different SCCI, which requires a representation that can be traversed in a controlled and interpretable way.

In this work, we propose an unsupervisedly learned skin-colour disentangling framework that treats SCCI as the outcome of multiple entangled factors. Inspired by recent applications [14, 26] of the information bottleneck principle [22] to low-level vision feature disentanglement, we adapt [26]’s disentanglement-by-compression framework and compress skin colour into an organized latent space whose entries are approximately independent and individually adjustable. Faithful disentanglement of SCCI possesses unique challenges compared to [26]. First, to avoid shortcuts that leak skin-darkness/shading information and obstruct the representation learning, we introduce a randomized, mostly monotonic decolorization mapping that generates the colourless input required by the framework. Second, to suppress the unintended colour shift of localized patterns (e.g., ink marks, scars), we propose a geometry-aligned post-processing step to selectively reject unreliable changes. The ability of faithful controlled manipulation in a latent space enables visualizing the lesion under a different SCCI.

After training, SCCI can be transferred directly from source images to targets (fig. 1(a)), and individual latent entries can be adjusted independently (fig. 1(b)). The proposed method improves health equity by adjusting the training set for different populations. Dataset-level augmentation can be performed by sampling from a target distribution (fig. 1(c)), and data normalization can be performed

by setting the colour embedding to a standard value. Physically meaningful trajectories can be identified and traversed in the latent space (fig. 2) for educational visualization. In summary, the main contributions of this work are:

- An unsupervised model that answers the counterfactual “*What would the skin condition look like if the skin colour captured in the image were different?*” via latent-space manipulation and transfer, enabled by the randomized decolourization module we propose.
- A geometry-aligned post-processing step that corrects unintended colour shifts, improving faithfulness of the intended appearance transformation.
- Augmentation and normalization strategies for downstream tasks using the proposed SCCI modelling method.

2 Methodology

2.1 Model Architecture for Perceived Skin Colour Modelling

Since perceived skin colour captured in images (SCCI) is influenced by multiple entangled factors (e.g., skin tone, illumination, and camera settings), we model it with a disentanglement-by-compression framework adapted from [26] (retaining most of its architecture). As shown in fig. 3(a), the colour encoding network e maps an input image $\mathbf{y}$ to a low-bitrate colour embedding $\mathbf{e} = e(\mathbf{y})$. A colour synthesis network f then reconstructs $\hat{\mathbf{y}} = f(\mathbf{x}, \mathbf{e})$ from $\mathbf{e}$ and a colourless image $\mathbf{x}$ that preserves geometry but suppresses colour from $\mathbf{y}$. The loss functions constraints $\mathbf{e}$ to contain just-enough information for faithful reconstruction [26], hence, encouraging e to retain primarily global SCCI information. At the end of training, only a few entries in $\mathbf{e}$ will be active (having non-negligible variance across the training set), allowing manual manipulation of image colour features. At test time, replacing $\mathbf{x}$ with another image enables colour transfer, while editing or sampling entries of $\mathbf{e}$ enables controlled colour manipulation.¹

Decolourization Unlike [26], which assumes a separately provided $\mathbf{x}$ for each $\mathbf{y}$, we propose a randomized decolourization module to generate $\mathbf{x}$ from $\mathbf{y}$. A fixed linear greyscale conversion can leak luminance cues correlated with skin darkness and shading, which weakens control over perceived darkness. We therefore use a randomized pixel-wise mapping $g_{\alpha} : [0, 1]^3 \rightarrow [0, 1]$ with fixed endpoints

$$g_{\alpha}(\mathbf{0}) = 0, \quad g_{\alpha}(\mathbf{1}) = 1, \quad (1)$$

where α is resampled for each training image in each epoch, and g_{α} is chosen to be mostly increasing in each channel. In practice, a linear combination of monotonic quadratic terms is used to balance simplicity and diversity in concavity:

$$g_{\alpha}(\mathbf{c}) := \sum_{i=1}^3 \sum_{j=i}^3 \alpha_{ij} c_i c_j + \sum_{i=1}^3 \sum_{j=i}^3 \beta_{ij} (1 - (1 - c_i)(1 - c_j)), \quad (2)$$

¹ The code for the colour model is available at <https://github.com/yangwenbo99/disentangle-skin-colour>.

where $\mathbf{c} = (c_1, c_2, c_3)$ is the RGB vector and $\boldsymbol{\alpha} = (\alpha_{ij}, \beta_{ij})_{i \leq j}$ is sampled such that $\sum_{i \leq j} (\alpha_{ij} + \beta_{ij}) = 1$ and $\alpha_{ij}, \beta_{ij} \geq -\epsilon$ for a small ϵ that is chosen to balance flexibility and monotonicity.

2.2 Post-Processing for Colour Correction

Because the model is designed to capture SCCI (a global feature across the entire image), localized colour patterns (e.g., ink marks, scars, irritations) may undergo undesirable colour shifts during transfer or sampling. We correct the colour of such regions by defining a post-processing operator P that suppresses changes at pixels that were already hard to reconstruct. Let $\hat{\mathbf{y}} = f(\mathbf{x}, \mathbf{e})$ with $\mathbf{e} = e(\mathbf{y})$, and let $\hat{\mathbf{y}}' = f(\mathbf{x}, \mathbf{e}')$ be the manipulated output, where $\mathbf{e}'$ is obtained from another image, sampling, or manual adjustment. We define a rejection weight $\mathbf{w}$ by

$$w_{ij} := h(d_{ij}), \quad \text{where } d_{ij} := \frac{1}{1 - \tau} \max(\|\mathbf{y}_{ij} - \hat{\mathbf{y}}_{ij}\| - \tau, 0), \quad (3)$$

τ is a small tolerance threshold, and $h : [0, 1] \rightarrow [0, 1]$ is monotone increasing. The post-processing operator P only accepts the change proposed by f when the reconstruction error d_{ij} is small:

$$P(\hat{\mathbf{y}}') := \mathbf{y} + (1 - \mathbf{w}) \odot (\hat{\mathbf{y}}' - \mathbf{y}). \quad (4)$$

2.3 Data Augmentation and Normalization for Downstream Tasks

A typical downstream ML model is trained on skin images with distribution S and tested on T . Labels for S are always available, whereas labels for T are not during training. In some scenarios, however, unlabelled examples from T are available in advance (e.g. when a clinical site has collected some patient images but diagnosis is not yet available). It is well known that downstream performance can degrade when the SCCI distributions of S and T mismatch.

Training Data Augmentation In physician education, exposure to skin colours in the population that they will serve is essential for building an accurate mental model. Similarly, a ML model can benefit from a training set augmented to match the SCCI distribution of a target cohort:

$$S_{\text{aug}} := \left\{ P \circ f \left(g_{\boldsymbol{\alpha}}(\mathbf{y}^{(i)}), \mathbf{e}^{(i,k)} \right) : \mathbf{y}^{(i)} \sim S, \mathbf{e}^{(i,k)} \sim e(T) \right\} \quad (5)$$

where f is trained on $S \cup T$ without labels and $\mathbf{e}^{(i,k)} \sim e(T)$ is sampled from the embedding distribution of the test set. Since it is needed to model the SCCI distribution of T , some unlabelled images from T are needed during training. We note that the transductive methodology has similarity nature of the usage of document embeddings in unsupervised document classification, as both attempt to construct a latent space to represent the test set's distribution [11, p. 422]. This mode suits deployment settings where unlabelled target-domain images are available before testing while diagnostic labels remain unavailable. This procedure only modifies the training set, with no change needed to the test-time procedure, and can be used to generate training examples for physician education.

Colour Normalization In contrast, colour normalization is more ML-centric. It maps every image to a “standard” colour, reducing variability across devices and acquisition conditions in both training and evaluation:

$$\begin{aligned} S_{\text{norm}} &:= \left\{ P \circ f \left(g_{\alpha}(\mathbf{y}^{(i)}), \bar{\mathbf{e}} \right) : \mathbf{y}^{(i)} \sim S \right\}, \text{ and} \\ T_{\text{norm}} &:= \left\{ P \circ f \left(g_{\alpha}(\mathbf{y}^{(i)}), \bar{\mathbf{e}} \right) : \mathbf{y}^{(i)} \sim T \right\} \end{aligned} \quad (6)$$

where $\bar{\mathbf{e}}$ is the mean colour embedding of the training set used as the “standard”. Unlike augmentation, normalization supports online use (at the cost of running f at test time of downstream tasks): each test image is processed independently without knowledge of T ’s distribution (and without needing examples from T during training), and f can be trained on S alone.

3 Experiments and Results

Each model is trained on an H100 GPU for 50 epochs using the same loss functions, training procedure, and optimization settings as in [26], except that we set $\lambda_{\text{bpp_g}} = 0.003$, $\lambda_{\text{diver}} = 0.05$, and $\lambda_{\text{color}} = 0.1$. Different downstream tasks may use different training sets. For qualitative experiments, we use the augmentation model trained on the training and test sets in [17] and ISIC-2020 [19] (both without labels). For post-processing, we set $\tau = 0.1$ and $h(d) = d^2$.

3.1 Qualitative Results of Skin Colour Matching and Sampling

We selected several challenging examples from SkinCap [28] and ISIC-2020 [19] as test cases and generated results by transferring the SCCI between them. As can be seen in fig. 1(a), our model effectively transfers the SCCI without altering the skin structure or lesion appearance.

By computing the variance of each entry in the 256-dimensional latent vector $\mathbf{e}$, we identified a small number of active entries. The results of individually varying the first five active entries are shown in fig. 1(b), each shows a semi-interpretable physical factor. In addition, we independently sample each entry of $\mathbf{e}$ from the empirical marginal distribution of the SkinCap test set to augment a training image. The results are shown in fig. 1(c). The figure confirms that the model produces diverse yet realistic skin colour variations.

3.2 Exploration of the Latent Space

Realistic physical meaning can be observed in the latent space. We captured a sequence of trajectory-building sample photos (TBSPs) corresponding to three physical factors: (a) skin blood perfusion (via temperature change), (b) camera white balance settings, and (c) the colour temperature of lights. For each factor, the TBSPs are fed into the colour encoding network to produce a sequence of latent vectors, which are connected in a piecewise-linear manner to form a sample

Table 1. Accuracy of lesion malignancy classification models trained on different augmented or normalized training sets. Results of existing methods are obtained from the original authors.

Index	Aug./Norm. Method	(Modification)	Acc. (std)
Baseline and Prior Arts			
1	(Baseline) [17]		0.561 (0.020)
2	Augmentation w/ Style Transfer [17]		0.761 (0.018)
Proposed Augmentation and Normalization Methods			
3	Proposed Augmentation (eq. (5))		0.772 (0.014)
4	Proposed Normalization (eq. (6))		0.764 (0.036)
Ablation Studies			
5	Proposed Augmentation (eq. (5))	(w/ Flow Sampler)	0.765 (0.028)
6	Proposed Augmentation (eq. (5))	(w/ Independent Sampler)	0.759 (0.014)
7	Proposed Normalization (eq. (6))	(w/o post-processing)	0.727 (0.009)
8	Proposed Normalization (eq. (6))	(w/ naive greyscale conversion)	0.722 (0.015)




Fig. 4. Qualitative ablation study results of transferring perceived skin colours. Rows and columns are defined in the same way as in fig. 1(a).

trajectory (STra) in the latent space. To visualise the effect, we randomly select images and vary their colour latent e along a curve parallel to the STra. Results from one test image are shown in fig. 2(a,b,c) (animated examples can be found in the supplementary materials). The PCA plot of the STra for camera white balance settings (fig. 2(d)) traces a smooth curve, indicating the model has learned a continuous representation of this physical factor.

3.3 Dataset Augmentation and Colour Normalization for Lesion Classification

Our method is a general framework applicable to various downstream tasks. Due to space constraints, we present only the results of training a lesion malignancy classification model. Note that the goal of these experiments is not to propose a state-of-the-art classifier, but to explore the potential of our augmentation and normalization methods. While an evaluation by a large panel of dermatologists would be valuable, the number of examples that can be evaluated would still be limited. Therefore, we use a diagnostic classification task as a proxy to evaluate the effectiveness of the proposed methods in clinical applications, following the protocol established in [17]. While other GAN-based augmentation methods have been proposed in the literature, their objectives differ from ours (their works

focus on modifying lesion attributes such as size, type and shape, while our method focuses on modifying the captured skin colour), and their evaluation protocols differ from [17], making direct comparison difficult.

For fair comparison, we obtained [17]’s training set (parts of DermNet NZ [5] and ISIC-2018 JID editorial images [12]) and test set (parts of Dermatology Atlas [15]) directly from the authors. For each augmented or normalized training set, we train a ResNet-50 [9] model for lesion malignancy classification using Adam with a learning rate of 10^{-3} , repeating each experiment ten times and reporting the mean and standard deviation of accuracy. The colour model for data augmentation is trained on unlabelled training and test sets for classification, while the colour normalization model is trained on the training set only (to improve representation quality, unlabelled ISIC-2020 [19] is added to the set). The models trained on the augmented (row 3 in table 1) and normalized (row 4) sets outperform the baseline (row 1) by a large margin. They also perform favorably compared to the prior augmentation method [17], while providing significantly more interpretability in the representation space.

3.4 Ablation Study

Effects of Post-Processing Certain non-skin features (e.g., scars and ink marks) that are faithfully preserved after post-processing (fig. 1(a)) are absent in the direct network output (fig. 4(a)). When post-processing is removed from the normalization pipeline, the downstream classifier performs significantly worse than the full method (row 7 vs. row 4 in table 1). We also computed the average LPIPS score [27] between the normalized and original images (after masking out non-lesion regions using masks provided by [18]), and found that removing post-processing worsens the score from 0.011 to 0.045.

Effects of Decolourization Methods The naive alternative to the proposed decolourization method is the default greyscale conversion method in PyTorch [2]. In contrast to fig. 1(a), the model trained with the naive method’s outputs produces results (fig. 4(b)) whose skin darkness resembles that of the target image (which should have provided lesion structure) rather than the source image (which should have provided SCCI). This indicates the naive conversion is not capable of enabling faithful skin colour transfer and can create significant obstacle to provide diverse augmentation results for education.

Effects of Latent Sampling Methods Under ideal conditions (a very large training set and a highly flexible model), the entries of $\mathbf{e}$ can be assumed independent; in practice, however, weak correlations between entries are observed. We therefore compare several latent sampling strategies for the augmentation method; results are shown at the bottom of table 1. We tested 3 methods for sampling the latent vector $\mathbf{e}$ from the distribution of the test set: (row 3) directly reusing latents from the test set (by random selection), (row 5) fitting a normalizing flow model [6], and (row 6) sampling each entry independently from the empirical marginal distribution. All methods with a relatively accurate modelling of the marginal distribution of the latent vector can achieve comparable performance, while direct random sampling performs the best.

4 Conclusion

In this work, we presented a disentanglement-by-compression framework that learns a structured, manipulable latent representation of skin colour captured in images (SCCI) from unlabelled dermatology data. A randomized decolourization mapping and a geometry-aligned post-processing step were proposed to improve faithfulness of the colour manipulation; ablation studies confirmed that both contribute measurably to downstream performance and perceptual fidelity. The learned latent space supports colour transfer, independent entry manipulation, and traversal along physically meaningful trajectories (e.g., blood perfusion, camera white balance), enabling training set augmentation, colour normalization, and educational visualization of skin conditions under varying SCCI. On a downstream benchmark, both augmentation and normalization strategies perform favourably against prior methods. This work is a step towards achieving equitable diagnosis, where our models can be reliably adapted across different clinical settings and populations.

Acknowledgments. This work is supported in part by NSERC Postdoctoral Award, NSERC Discovery Grant, Canada Research Chair Programs, Ontario Early Researcher Award, SHARCNET, Calcul Québec, and the Digital Research Alliance of Canada.

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

References

1. Aggarwal, P., Papay, F.A.: Artificial intelligence image recognition of melanoma and basal cell carcinoma in racially diverse populations. *Journal of Dermatological Treatment* **33**(4), 2257–2262 (2022). <https://doi.org/10.1080/09546634.2021.1944970>
2. Ansel, J., Yang, E., He, H., et al.: PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In: 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS ’24). ACM (Apr 2024). <https://doi.org/10.1145/3620665.3640366>
3. Benmalek, A., Cintas, C., Tadesse, G.A., Daneshjou, R., Varshney, K.R., Dalila, C.: Evaluating the impact of skin tone representation on out-of-distribution detection performance in dermatology. In: Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5 (2024). <https://doi.org/10.1109/ISBI56570.2024.10635847>
4. Cong, C., Liu, S., Di Ieva, A., Pagnucco, M., Berkovsky, S., Song, Y.: Semi-supervised adversarial learning for stain normalisation in histopathology images. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Es-sert, C. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 581–591. Springer International Publishing, Cham (2021)
5. DermNet NZ: DermNet NZ, <https://dermnetnz.org/>
6. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real nvp. arXiv preprint arXiv:1605.08803 (2016)

7. Fitzpatrick, T.B.: Soleil et peau. *J. Med. Esthet.* **2**, 33–34 (1975)
8. Ghorbani, A., Natarajan, V., Coz, D., Liu, Y.: DermGAN: Synthetic generation of clinical skin images with pathology. In: *Machine learning for health workshop*. pp. 155–170. PMLR (2020)
9. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1026–1034 (2015)
10. Howard, J.J., Sirotin, Y.B., Tipton, J.L., Vemury, A.R.: Reliability and validity of image-based and self-reported skin phenotype metrics. *IEEE Transactions on Biometrics, Behavior, and Identity Science* **3**(4), 550–560 (2021)
11. Incitti, F., Urli, F., Snidaro, L.: Beyond word embeddings: A survey. *Information Fusion* **89**, 418–436 (2023). <https://doi.org/10.1016/j.inffus.2022.08.024>
12. International Skin Image Collaboration: International skin image collaboration, <https://www.isic-archive.com/>
13. Iyatomi, H., Celebi, M.E., Schaefer, G., Tanaka, M.: Automated color calibration method for dermoscopy images. *Computerized Medical Imaging and Graphics* **35**(2), 89–98 (2011). <https://doi.org/https://doi.org/10.1016/j.compmedimag.2010.08.003>, <https://www.sciencedirect.com/science/article/pii/S0895611110000832>, advances in Skin Cancer Image Analysis
14. Li, X., Chen, Q., Peng, X., Yu, K., Chen, X., Lu, Y.: Bitrate-controlled diffusion for disentangling motion and content in video. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12904–12914 (2025)
15. Marques, B.P., Melo, C.C., Calheiros, D.B., Portugal, F.M., Junior, G.B.L., Marcolino, G.C., Amorim, H., Oliveira, I.C.D., Cruz, L.M.O., Filho, L.L.L., Botelho, M.C., Brandão, P.L.M., Prado, R., Pérez, R.H., Santana, S., Mendonça, T.A., Macedo, V.A.O.M., Nascimento, V.L.B., Inakanti, Y.: *Dermatology atlas*, <https://www.atlasdermatologico.com.br/contributors.jsf>
16. Monk, E.: The monk skin tone scale (2023). <https://doi.org/https://doi.org/10.31235/osf.io/pdf4c>
17. Rezk, E., Eltorki, M., El-Dakhakhni, W.: Improving skin color diversity in cancer detection: Deep learning approach. *JMIR Dermatology* **5**(3) (2022). <https://doi.org/https://doi.org/10.2196/39143>, <https://www.sciencedirect.com/science/article/pii/S2562095922000605>
18. Rezk, E., Eltorki, M., El-Dakhakhni, W.: Leveraging artificial intelligence to improve the diversity of dermatological skin color pathology: protocol for an algorithm development and validation study. *JMIR Research Protocols* **11**(3), e34896 (2022)
19. Rotemberg, V., Kurtansky, N., Betz-Stablein, B., Caffery, L., Chousakos, E., Codella, N., Combalia, M., Dusza, S., Guitera, P., Gutman, D., et al.: A patient-centric dataset of images and metadata for identifying melanomas using clinical context. *Scientific data* **8**(1), 34 (2021)
20. Salvi, M., Branciforti, F., Veronese, F., Zavattaro, E., Tarantino, V., Savoia, P., Meiburger, K.M.: DermoCC-GAN: A new approach for standardizing dermatological images using generative adversarial networks. *Computer Methods and Programs in Biomedicine* **225**, 107040 (2022). <https://doi.org/10.1016/j.cmpb.2022.107040>
21. Schaefer, G., Rajab, M.I., Emre Celebi, M., Iyatomi, H.: Colour and contrast enhancement for improved skin lesion segmentation. *Computerized Medical Imaging and Graphics* **35**(2), 99–104 (2011). <https://doi.org/10.1016/j.compmedimag.2010.08.004>, advances in Skin Cancer Image Analysis

22. Tishby, N., Pereira, F.C., Bialek, W.: The information bottleneck method. arXiv preprint physics/0004057 (2000)
23. Tosta, T.A.A., Freitas, A.D., de Faria, P.R., Neves, L.A., Martins, A.S., do Nascimento, M.Z.: A stain color normalization with robust dictionary learning for breast cancer histological images processing. *Biomedical Signal Processing and Control* **85**, 104978 (2023). <https://doi.org/10.1016/j.bspc.2023.104978>
24. Weir, V.R., Li, Y., Gillis, M.C., Kurtansky, N.R., Salvador, T., Halpern, A.C., Nelson, K.C., Lester, J.C., Rotemberg, V.: Evaluating skin tone scales for dermatologic dataset labeling: a prospective-comparative study. *npj Digital Medicine* **8**(1), 787 (2025)
25. Xu, C., Sun, Y., Zhang, Y., Liu, T., Wang, X., Hu, D., Huang, S., Li, J., Zhang, F., Li, G.: Stain normalization of histopathological images based on deep learning: A review. *Diagnostics* **15**(8), 1032 (2025)
26. Yang, W., Wang, Z., Wang, Z.: Towards a universal image degradation model via content-degradation disentanglement. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12966–12975 (2025)
27. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
28. Zhou, J., Sun, L., Xu, Y., Liu, W., Afvari, S., Han, Z., Song, J., Ji, Y., He, X., Gao, X.: Skincap: A multi-modal dermatology dataset annotated with rich medical captions (2024)