IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received February 18, 2021, accepted March 2, 2021, date of publication March 12, 2021, date of current version March 23, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3065701

Unsupervised Feature Elimination via Generative Adversarial Networks: Application to Hair Removal in Melanoma Classification

DAHYE KIM AND BYUNG-WOO HONG
School of Computer Science and Engineering, Chung-Ang University, Seoul 06974, South Korea

Corresponding author: Byung-Woo Hong (hong@image.cau.ac.kr)

This work was supported in part by the Chung-Ang University Graduate Research Scholarship in 2019, and in part by the National Research Foundation of Korea under Grant NRF-2017R1A2B4006023 and Grant NRF-2019K1A3A1A77074958.

ABSTRACT Eliminating the undesirable features is crucial to computer vision applications since undesirable features degrade the visibility of images. For that purpose, denoising, dehazing and deraining have been actively studied in both traditional model-based approaches and modern deep learning methods. However, elimination of hair in dermoscopic images has not received sufficient attention despite its significance and potential. Meanwhile, hair removal algorithms remain within the classical morphological methodologies, while only a few attempts apply the latest data-driven techniques. Hair is desired to be removed in dermoscopy applications because it causes undesired effects such as occlusions in lesion areas. However, removing hair is challenging because of its inherent complex structure and variations. In this paper, we propose a new unsupervised algorithm for hair removal and evaluate it on a real-world melanoma dataset. The proposed method eliminates hair from dermoscopic images by inducing a reconstructed distribution of images with hair to resemble a hairless distribution using generative adversarial learning. In the generative adversarial learning framework, hair features are characterized with a coarse-grained label simply via a binary classifier. At the same time, the important features of the lesions are preserved by minimizing L_1-norm reconstruction loss based on Laplace noise assumption. The qualitative evaluation of the hair-removed results show that the proposed algorithm is robust to hair variations, and the quantitative results demonstrate that applying our hair removal algorithm considerably improves the performance of melanoma classification, outperforming the benchmarks.

INDEX TERMS Deep learning application, dermoscopy, generative adversarial networks, hair removal, medical imaging, skin lesion classification, unsupervised learning.

I. INTRODUCTION

Undesirable features such as noise, haze and rain degrade the performance of computer vision algorithms since they reduce the visibility of images [1]. Feature elimination tasks aim to remove the undesirable features from images to get clear representations that are ideal for computer vision applications [2], [3]. Feature elimination applications like denoising, dehazing, and deraining have been actively studied. Classical approaches mainly rely on characterizing the statistical properties of undesirable features [4]–[22]. Progress in deep learning has led to data-driven solutions where fine-grained annotations are used as ground-truth [16], [23], [23]–[40].

The associate editor coordinating the review of this manuscript and approving it for publication was Donato Impedovo.

With the advent of a generative adversarial network (GAN) [41] and CycleGAN [42], data-driven unsupervised techniques where undesirable features are removed by imposing the distribution of clean images to the reconstructed output via adversarial loss have been developed [39], [43]–[46]. Hair is an undesirable factor that makes visual degradations and causes serious occlusions in dermoscopic images. However, removing hair from dermoscopic images has not been sufficiently investigated under the deep learning framework despite its importance and potentials in medical imaging. For example, the classification of melanoma, known as fatal skin cancer, is challenging due to the destructive interruption by hair on the skin. Hair disturbs the intrinsic intensities and geometric properties of the lesion regions causing detrimental effects on classification. This chronic problem raises the

need for hair removal pre-processes on dermoscopic images. Classical approaches often employ morphological methods to remove hair by defining the problem as a hair region detection and inpainting task [47]–[53]. They use pre-defined filters to detect hair regions and change the intensity values of the hair areas through interpolation. However, the conventional methods have difficulties in characterizing hair regions and finding suitable interpolation methods in images containing hairs with complex structures and large variations. Recently, the data-driven deep learning technique has been adopted to solve the hair removal problem in a supervised way with fine-grained hairless labels that are expensive to construct and often unavailable [54]. Meanwhile, GAN has also been applied to dermoscopic images. However, the application of GAN in dermoscopic images is limited to synthetic image generation for data augmentation [55]–[59] and segmentation improvement [60], [61]. In this paper, we introduce an unsupervised undesirable feature elimination algorithm and demonstrate it by applying it to hair removal. The proposed algorithm eliminates undesirable features from input images by imposing the distribution of the real clean images via GAN [41] in unpaired dataset. Under the GAN framework, the proposed algorithm characterizes the undesirable features to be removed with a coarse-grained label via a simple binary classifier, the discriminator of GAN. Simultaneously, prior supervised methods require fine-grained labels that include location, shape and other detailed properties of the feature. The proposed method combines image reconstruction and translation. The image reconstruction part is built upon the Laplace noise assumption, and it promotes the preservation of critical original properties of the images. The translation part uses GAN to induce the obtained reconstructed images to follow the distribution of clean images where the undesirable features are not observed. The two components of the algorithm are simultaneously optimized under the proposed integrated loss function for the undesirable feature elimination. The effectiveness of the proposed algorithm is tested on dermoscopic images for melanoma diagnosis. To our best knowledge, our method provides the first demonstration of a data-driven unsupervised hair elimination technique in dermoscopic imaging. In our proposed algorithm, hair, the undesirable feature, is characterized by the discriminator of GAN, which distinguishes between images with hair and without hair. The reconstruction part preserves the important original properties such as shape, location and color of the images' lesions. The hairless translation based on GAN enforces the distribution of the obtained reconstructed images to be similar to that of the hairless images. The proposed algorithm shows hair-removed results that are robust to variations of hair properties such as density, thickness and color as shown in Fig. 5, and Fig. 6. Furthermore, the impact of hair removal by the proposed method is evaluated on melanoma classification. The hair elimination algorithm is trained beforehand and employed to pre-process hair-containing images. We compare the classification performance on the original images, and the reconstructed images where the hair is removed by the benchmarks and our proposed technique. As demonstrated by Tab. 2 and 3, the experimental results show that removing hair from the input images by the proposed algorithm significantly improves the performance of state-of-the-art deep neural networks for classification [62]–[64]. The contributions of this paper are listed as follows:

- We present a new unsupervised data-driven approach for undesirable feature elimination.
- We propose an algorithm to characterize features with a coarse-grained label via a simple binary classifier.
- We propose a new image reconstruction algorithm incorporating the feature manipulation in the framework of generative adversarial networks.
- We present the first demonstration of an unsupervised deep learning approach for hair removal from dermoscopic images.
- We improve an image representation that can enhance the accuracy in the classification of melanoma lesions in skin images.

II. RELATED WORK

In feature elimination, denoising, dehazing and deraining problems have been the main research topics. The general approaches used in conventional methods utilizes characterizing statistical properties of the features to be removed [4]–[22]. More recently, deep learning has been employed for feature elimination tasks and fine-grained annotations are used as ground-truth for supervised learning [16], [23], [23]–[40]. Some supervised approaches combine GAN for reducing artifacts [36], [40] and augmenting paired training data with undesirable features and ground-truth clean images [37]–[39]. The development of CycleGAN [42], which is originally proposed for unpaired translation between two different image domains based on cycle-consistency loss, has inspired unsupervised techniques in feature elimination tasks [45], [46]. Unsupervised feature elimination based on CycleGAN is further improved with a combination of perceptual loss used to preserve the original contents of images [44] for deraining. The most similar approach to our proposed model is given by UNet+L2, the recent denoising method for X-Ray CT images [43] where the L_2-norm is used for reconstruction loss under the Gaussian noise assumption and GAN is adopted to impose the clean image distribution. Our proposed model is also based on the GAN framework but we employ the L_1-norm that is known to be more preferred because of its anisotropic and sparse properties, assuming that the additive noise in images follows Laplace distribution. Moreover, we test our algorithm on the real-world dermoscopic data while only a synthetic noise denoising case is evaluated in [43]. Classical approaches have considered hair removal from images as a hair detection and occlusion inpainting problem [47]–[49] and used filter-based methods to solve it. They use morphological closing operations to obtain binary masks of hair and attempt to inpaint hair regions by traditional interpolation methods [47], PDE (partial differential equation) [48] or modified coherence

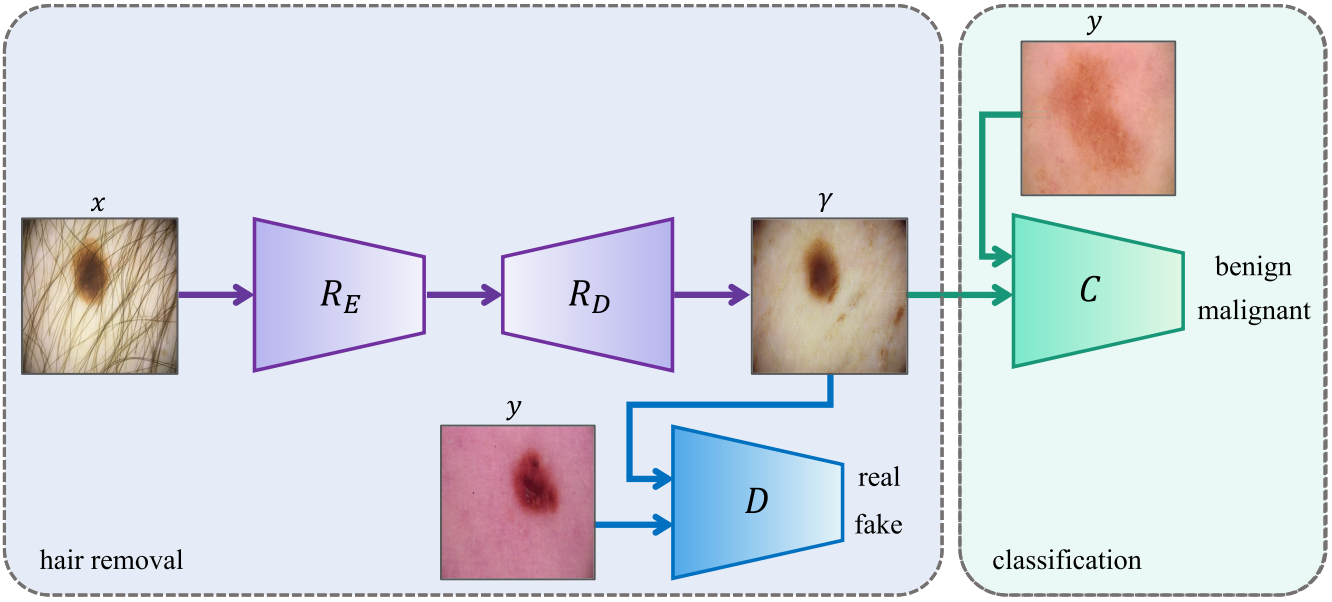

FIGURE 1. Schematic illustration of the proposed neural network architecture. Neural networks for hair removal consist of the reconstructor *R* that translates the input image with hair *x* to the hairless image γ and the discriminator *D* that classifies the real hairless image *y* and the reconstructed hairless image γ. In classification, the image with hair *x* is pre-processed by the pre-trained hair removal reconstructor *R* and the reconstructed hairless result γ is passed to the melanoma classifier *C*, while the real hairless image *y* is directly fed to the *C* without pre-processing for hair elimination.

transport [49]. More recently, multi-scale curvilinear matched filters are applied in grayscale images to detect hairs in images and region growing algorithms, and the linear discriminant analysis (LDA) are used to inpaint the detected hair regions in [50]. An adaptive canny edge detector is leveraged to detect hair regions, and the regions are inpainted using a wavelet-based multi-resolution coherence transport [51]. Modified morphological closing approaches are applied for hair detection and inpainting both in [53]. Modern data-driven approaches for hair removal adopt deep neural network auto-encoders and use supervised training with fine-grained annotations that contain specific information of hair such as location, shape and color [54]. The auto-encoder-based method eliminates hair from dermoscopic images by minimizing the discrepancy between the auto-encoder's resulting output and the fine-grained ground-truth [54]. However, the supervised hair removal approach requires pairs of hair-containing images and their ground-truth hair-removed images that are expensive to construct and often unavailable. On the contrary, our proposed method leverages the deep neural auto-encoder network trained in an unsupervised manner where the ground-truth hair-removed images are not used, and the GAN framework is employed to learn the distribution of hairless data. GAN is a modern deep learning technique where a generative model generator and a discriminative model discriminator are trained in an adversarial way based on Jensen-Shannon divergence [41]. The generator produces a similar distribution to the given real data from random vectors while the discriminator distinguishes between the real data and the generated fake data. GAN framework can be integrated with various computer vision tasks like image translation, domain adaptation, super-resolution and image generation [42], [65]–[67]. GAN has been applied in medical domain vision problems generally for synthetic data generation [68]–[70], domain adaptation [71], [72] and segmentation [73], [74]. Particularly, dermoscopy application has adopted GAN for the realistic synthetic image generation for data augmentation [55]–[59] and enhancing segmentation performance [60], [61]. However, to the best of our knowledge, using the GAN framework for undesirable feature removal in dermoscopy is still an open question, and we propose the first data-driven unsupervised approach where the GAN framework is leveraged to eliminate hair from dermoscopic images in an unsupervised way.

III. METHOD

Our objective is to learn mapping function R from hair domain X to hairless domain Y, $R: X \rightarrow Y$, while preserving other properties such as location and shape of lesions and intensities given training samples $x_i{}_{i=1}^{N}$ where $x_i \in X$ and $y_j{}_{j=1}^{M}$ where $y_j \in Y$. The proposed hair removal process consists of image reconstruction and hair-to-hairless translation.

A. RECONSTRUCTION

The problem of interest is to obtain an optimal reconstruction γ given an observation x by maximizing the posterior probability $p(\gamma \mid x)$ that is proportional to the product of the likelihood probability $p(x \mid \gamma)$ and the prior probability $p(R)$ due to the Bayes rule as follows:

$$p(\gamma \mid x) \propto p(x \mid \gamma)p(\gamma). \tag{1}$$

The likelihood probability is designed to measure a dissimilarity between an observation x and its desired reconstruction γ following an additive image model as given by:

$$\gamma = x + \eta, \tag{2}$$

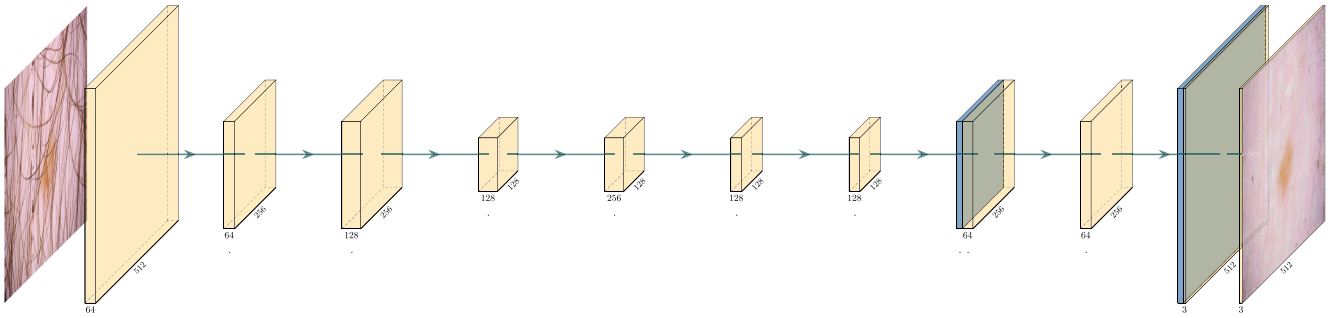

FIGURE 2. Illustration of the proposed network of the reconstruction model R. The blue arrows show the flow of activation. The yellow blocks represent the output activation of convolutional layers with 3 × 3 convolutional kernels and the blue blocks stand for the results of nearest upsampling.

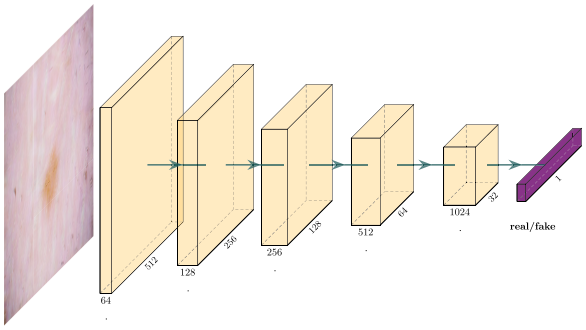

FIGURE 3. Illustration of the proposed network D of the discriminator. The blue arrows show the flow of activation. The yellow blocks represent the output activation of convolutional layers with 3 × 3 convolutional kernels.

where the noise process η is assumed to follow a Laplace distribution $L(0, b)$ with mean 0 and a diversity parameter $b > 0$. Then, we have the likelihood probability density as follows:

$$p(x \mid \gamma) \propto \exp\left(-\|x - \gamma\|_1\right), \tag{3}$$

where the L_1-norm between the observation x and model R is known to be effective in characterizing features of the model in an anisotropic way due to the sparsity assumption. The prior probability $p(\gamma)$ is designed to impose a constraint penalizing the proximity to the hair features.

1) RECONSTRUCTION LOSS

The reconstruction γ from an observation x is obtained by minimizing the following loss function:

$$\mathcal{L}_{\text{recon}}(\gamma, X, Y) = \mathbb{E}_{x\sim p_{\text{hair}}(x)}[\|\gamma - x\|_1], \tag{4}$$

which is derived by taking the negative logarithm of the likelihood probability in Eq. (3). The loss function $\mathcal{L}_{\text{recon}}$ measures the discrepancy between the observation x and its desired reconstruction γ in terms of the mean absolute error (MAE), thus the reconstruction is constrained to be similar to the observation in appearance.

B. HAIR TO HAIRLESS TRANSLATION

The elimination of undesired features associated with hairs can be achieved by the transformation of one domain X with hairs to the other Y without hairs. The domain transformation is developed in the generative adversarial network framework where the probability density function associated with the domain X is transformed to the probability density function associated with the domain Y, which is imposed by the extended prior probability $p(\gamma, y)$ with an additional random variable y representing an element in the domain Y without hairs. The joint prior probability $p(\gamma, y)$ is given by the product of the conditional probability $p(\gamma, y)$ and the probability $p(y)$ as follows:

$$p(\gamma, y) = p(\gamma \mid y)\, p(y). \tag{5}$$

The conditional probability $p(\gamma \mid y)$ measures how likely the reconstruction γ is drawn from the distribution Y representing the hairlessness.

1) ADVERSARIAL LOSS

We assume that the conditional probability $p(\gamma \mid y)$ in Eq. (5) is constructed based on Jensen-Shannon divergence JSD. The Jensen-Shannon divergence JSD provides a discrepancy measure between the probability distribution of the reconstructed results $q(\gamma)$ and the probability distribution of the hairless images $p(y)$ as the following:

$$-\log p(\gamma \mid y) \propto \text{JSD}(q(\gamma) \parallel p(y)). \tag{6}$$

Imposing the hairlessness to the resulting reconstructed images is corresponding to translating the domain of the hair images to the hairless domain. The translation from hair to hairless can be obtained by minimizing Jensen-Shannon divergence JSD in Eq. (6). Minimizing Jensen-Shannon divergence JSD can be solved by GAN mechanism [41]. Under the GAN framework, the objective is formulated by the minimax problem as follows:

$$\min_R \max_D \left(\mathbb{E}_{y\sim p_{\text{hairless}}(y)}[\log D(y)] + \mathbb{E}_{x\sim p_{\text{hair}}(x)}[\log(1 - D(\gamma_x))]\right), \tag{7}$$

where γ_x denotes a reconstruction associated with $x \sim p_{\text{hair}}$, D denote a discriminator to discriminate the real hairless image y and the reconstructed fake hairless image γ and $y \sim p_{\text{hairless}}(y)$ and $x \sim p_{\text{hair}}(x)$ stand for data distribution of hair data and hairless data, respectively. Based on the GAN objective in Eq. (7), we employ the non-saturating version of

Algorithm 1 Hair Removal via GAN. Default Values: $k = 10$, $\lambda = 100$

Require: Mini-batch size m, number of training iterations T
Require: Initial discriminator parameters w_0, initial reconstructor parameters θ_0
for $t = 1, \ldots, T$ **do**
  Sample mini-batch of hairless images $\{y^{(1)}, \ldots, y^{(m)}\} \sim p_{\text{hairless}}$.
  Sample mini-batch of images with hair $\{x^{(1)}, \ldots, x^{(m)}\} \sim p_{\text{hair}}$.
  $\gamma_x \leftarrow R(x)$
  $L_{\text{adv}}(w; y, x) = \frac{1}{m}\sum_{i=1}^{m}\left[\log D(y^{(i)}) - \log D(\gamma_{x^{(i)}}) - \frac{k}{2}\|\nabla D(y^{(i)})\|_2^2\right]$
  $w \leftarrow w + \nabla_w L_{\text{adv}}(w; y, x)$ # gradient ascent by maximizing L_{adv} for D_w
  $L_{\text{adv}}(\theta; x) = \frac{1}{m}\sum_{i=1}^{m}\left[-\log D(\gamma_{x^{(i)}})\right]$
  $\mathcal{L}_{\text{recon}}(\theta; x) = \frac{1}{m}\sum_{i=1}^{m}\left[\|\gamma_{x^{(i)}} - x^{(i)}\|_1\right]$
  $\theta \leftarrow \theta - \nabla_\theta\left(L_{\text{adv}}(\theta; x) + \mathcal{L}_{\text{recon}}(\theta; x)\right)$ # gradient descent by minimizing L_{adv} and $\mathcal{L}_{\text{recon}}$ for R_θ
end for

vanila GAN loss with a regularizer to penalize the gradients of the discriminator D on the real hairless image y [75].

$$\begin{aligned} &L_{\text{adv}}(w, \theta; Y, \gamma) \\ &= \mathbb{E}_{y\sim p_{\text{hairless}}(y)}[\log D(y)] \\ &\quad - \mathbb{E}_{x\sim p_{\text{hair}}(x)}[\log D(\gamma_x)] - \frac{k}{2}\mathbb{E}_{y\sim p_{\text{hairless}}(y)}[\|\nabla D(y)\|_2^2] \end{aligned} \tag{8}$$

where w and θ denote the trainable model parameters of the discriminator D and the reconstructor R which outputs γ_x, and k is a positive integer value to decide the degree of the regularization.

2) FULL OBJECTIVE FOR HAIR REMOVAL

The reconstruction loss is simultaneously applied with the adversarial loss. To control the level of imposition of hairlessness, we add λ as the pre-defined coefficient for the adversarial loss L_{adv}.

$$L_{\text{hair\_removal}}(\gamma, D, X, Y) = L_{\text{recon}}(\gamma, X, Y) + \lambda L_{\text{adv}}(w, \theta; Y, \gamma) \tag{9}$$

C. NEURAL NETWORK ARCHITECTURES

The proposed algorithm includes three separate neural networks; the reconstructor R, the discriminator D and the melanoma classifier C. The full architecture of the proposed method is trained in two sequential stages consisting of optimization of the reconstructor R and the discriminator D for removing hair, and optimization of the melanoma classifier C. The schematic illustration of the architectures is shown in Fig. 1. In the first stage, the proposed hair removal algorithm is optimized. The reconstructor R and the discriminator D are simultaneously trained in an end-to-end way. Let $R(x; \theta)$ be an auto-encoder for the image reconstruction with trainable model parameters θ. The R_E and R_D denote the encoder part and the decoder part of the reconstruction auto-encoder. The reconstructed image γ is fed into the discriminator D with the hairless data y. The discriminator D classifies the fake data γ and the real data y while the reconstruction γ tries to generate realistic hairless image from the hair input to fool the discriminator. This adversarial training scheme promotes the reconstructed results γ to follow the data distribution of real hairless data set Y, effectively removing hair from the input hair data x. The detailed neural network architectures of the reconstructor R and the discriminator D are shown in Fig. 2 and 3, respectively. In the second stage, the classifier is trained for the melanoma classification using the cross-entropy loss function and the parameters θ of the reconstructor R are frozen so that they are not affected by the optimization of the melanoma classifier C. If the given image is the hair image $x \in X$, the given x is first passed to the trained reconstructor R and the resulting reconstruction image γ is used as the input for the melanoma classifier C. On the other hand, if the given image is the hairless image $y \in Y$, the given y is directly used as the input for the melanoma classifier C and the trained reconstruction γ is not employed in this case.

While the adversarial loss induces the reconstruction output γ to follow the distribution of hairless data, the reconstruction loss attempts to preserve original properties of x. As a result of the simultaneous training with the reconstruction loss and the adversarial loss, we obtain hair-removed outputs where the other original properties of hair inputs remain. The training process for the proposed hair removal algorithm is described in Algorithm 1.

IV. RESULTS

The proposed algorithm is evaluated on the real-world dermoscopic dataset, The ISIC 2020 Challenge Dataset [76]. Due to the lack of prior data-driven unsupervised techniques that aim to eliminate hair from dermoscopic images, CycleGAN [42] which has been actively used for unsupervised feature elimination for denosing [39], dehazing [45], [46] and deraining [44], [77] as discussed in Sec. II and UNet+L2 [43] which is the most similar approach to ours are used as benchmarks. Although UNet+L2 is originally evaluated on denoising X-ray CT images, it is applicable to hair removal since its fundamental target problem is eliminating undesirable features from images in a data-driven unsupervised way. We compare qualitative hair removal results of our proposed algorithm and the benchmarks. Additionally, we quantitatively compare the classification performances on the original images, benchmark hair-removed images and our hair-removed images.

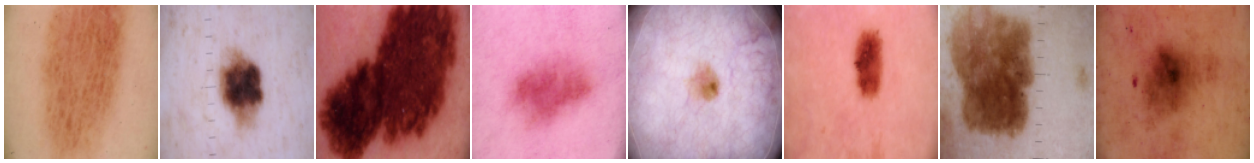

FIGURE 4. Examples of the skin lesion images without hair in the ISIC 2020 Challenge Dataset. The first four columns show examples of benign and the last four columns show examples of malignant.

The proposed algorithm is implemented following the details described in Fig. 2 and Fig. 3.

A. DATASET

The ISIC 2020 Challenge Dataset [76] that consists of 33,126 pairs of dermoscopic images of skin lesions obtained from over 2,000 patients and metadata is used in our experiments for hair removal algorithm and melanoma classification. The metadata includes binary labels of melanoma classification, but not whether each data contains hair or not.

B. HAIR REMOVAL

1) DATA

Since the GAN framework requires real data distribution that the generator wishes to imitate and our proposed algorithm needs hairless data to be used as the real data, we constructed hair labels for the ISIC 2020 Challenge Dataset. We selected the first 7,996 data and qualitatively evaluated whether each data has hair or not. Based on the evaluation, we added hair labels to the metadata. The constructed dataset with hair labels is composed of 5,007 hairless data and 2,989 hair-containing data. Note that assigning hair labels to the data is fundamentally different from building ground-truth annotations of hair-removed images. Using hairless distribution based on the GAN framework corresponds to imposing prior information to the generative model and training for hair removal is performed in the unsupervised manner whereas employing ground-truth of hair-removed images is a discriminative approach for hair removal with supervision with fine-grained annotations. For the hair removal algorithm, we trained the reconstructor R and the distriminator D by the whole data of the constructed dataset with hair labels. The hairless data samples of the constructed dataset are shown in Fig. 4 and the samples of images with hairs are given in Fig. 5 and Fig. 6. When we trained the reconstructor R and the discriminator D, we used the entire data of the constructed dataset with hair labels and resize its resolution to 512×512.

2) EVALUATION

In this section, we provide a qualitative evaluation of the proposed algorithm in the elimination of hairs presented in the skin images. The algorithm aims to identify the features characterizing hairs in the images, and subsequently get rid of them in an unsupervised way. For experiments of hair removal on the presented algorithm and the benchmarks, we use mini-batch sizes of 8 and train the reconstructor R and the discriminator D with learning rate 1e-04 for 186,500 iterations. We set k in Eq. (8) and λ in Eq. (9) as 10 and 100, respectively. For experiments of CycleGAN, we adopted the same structure of discriminator and reconstructor as our shown in Fig. 3 and Fig. 3. When experiment UNet+L2, we used UNet based autoencoder as presented in the paper [43] and we adopted the same discriminator as ours shown in Fig. 3 since the discriminator in the paper is designed for training with patches of an image however hair removal task requires training with a whole image. The examples of the original images and the hair-removed results by the benchmarks and our proposed algorithm are presented in Fig. 5 and Fig. 6 with benign and malignant lesions, respectively. In Fig. 5 and Fig. 6, the first row shows the original images that contain hair and the second and third rows present the reconstruction results after the elimination of hairs by benchmarks and our proposed algorithm, respectively. The qualitative results demonstrate that our proposed model shows more robust results on images with various types of hair than the benchmarks. The proposed algorithm and the benchmarks are compared on images of various hair densities. For example, the original images with sparse, moderate and dense hairs and their hair-removed results are presented in the first, second and third columns of Fig. 5 and Fig. 6, respectively. It is obvious that the occlusion that is desired to be recovered due to the hairs is proportional to the density of hairs, and thus it is considered to be more difficult to reconstruct images with dense hairs. Despite the difficulty, our proposed algorithm shows more robust results on images with different hair densities than the benchmarks. The examples in the third columns of Fig. 5 and Fig. 6 indicate that in images with excessively dense hairs, using our proposed model better preserves the original characteristics of the lesions whereas using the benchmarks damages the original properties of lesions. At the same time, images with various hair lengths are used for evaluations of the hair removal algorithms. For example, The forth columns of Fig. 5 and Fig. 6 present the original images with short hairs and their hair-removed images while the other columns display images with relatively longer hairs. Additionally, the hair elimination algorithms are tested on images with hairs of different thicknesses. For example, the fifth and sixth columns of Fig. 5 and Fig. 6 show the original images with thin hairs and thick hairs and their hair-removed results, respectively. Furthermore, images with different skin and lesion colors are used to evaluate the hair removal algorithms. For example, the original images of red

FIGURE 5. Examples of the benign skin lesion images with hair in the ISIC 2020 Challenge Dataset. Samples of original images with hair, hair-removed results by the benchmarks and hair-removed results by our proposed algorithm are shown from top to bottom row.

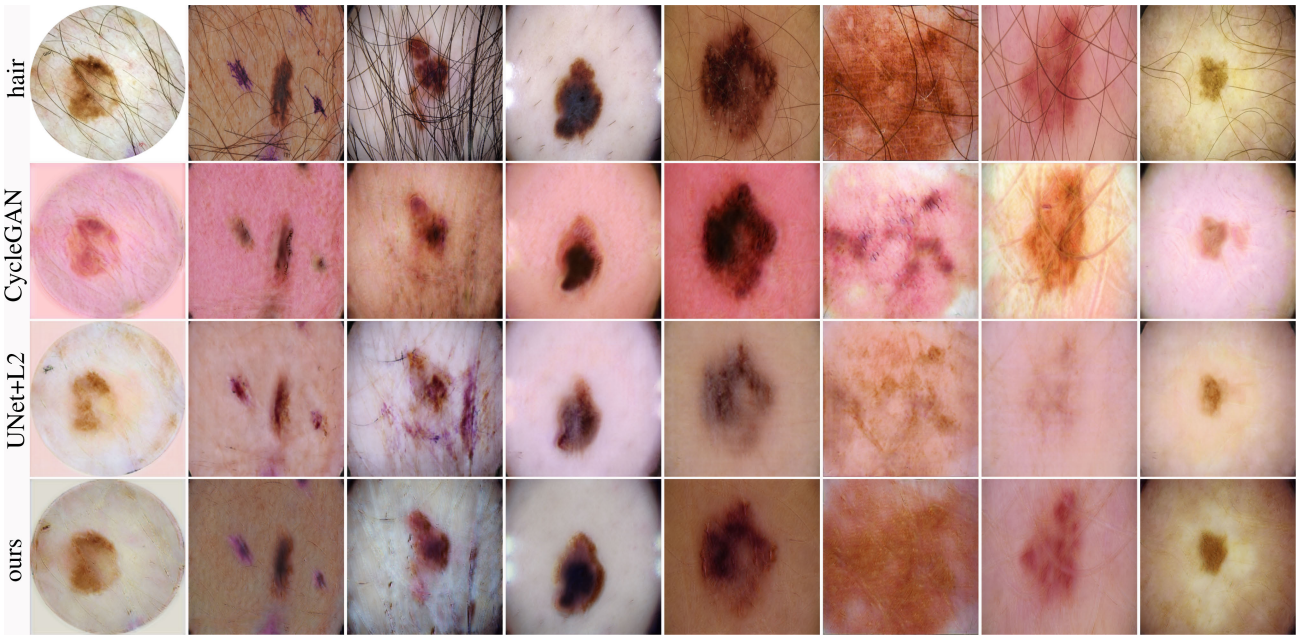

FIGURE 6. Examples of the malignant skin lesion images with hair in the ISIC 2020 Challenge Dataset. Samples of original images with hair, hair-removed results by the benchmarks and hair-removed results by our proposed algorithm are shown from top to bottom row.

and yellow skins and their hair-removed results are shown in the seventh and eighth columns of Fig. 5 and Fig. 6. The comparisons of hair-removed results between the benchmarks and our proposed algorithm show that our proposed algorithm is more effective in the preservation of the original shape of lesions and colors of skin and shows more robust hair elimination results on variations of hair than the benchmarks. More specifically, CycleGAN results in more distorted colors and shape compared to UNet+L2 and ours. For example, images in the sixth column of Fig. 5 and the seventh column of Fig. 6 show that CycleGAN converts yellow to red and significantly deforms the original shape of a lesion, while UNet+L2 and ours preserves the original intensity values and shapes. This indicates that training a native GAN combined with a reconstruction loss is better for preserving the original properties than training GAN in a cyclic way with the cycle-consistency

TABLE 1. Classification performances of EfficientNet-B4 on test set of the entire dataset where the distribution of labels for melanoma is severely imbalanced with only 1.7259% of malignant samples. The all column shows the accuracy on all test data and the hair column shows the accuracy only on test data with hair.

input	all	hair
original	0.9687	0.9565
CycleGAN	0.9800	0.9891
UNet+L2	0.9738	**0.9946**
ours	**0.9813**	**0.9946**

loss. Meanwhile, our method is even better in preserving the original properties than UNet+L2. For instance, the second column of Fig. 5 and the seventh and eighth columns of Fig. 6 show that our method better preserves the original colors and shapes of lesions. This implies that adopting a network without skip-connection as a reconstructor and using L_1-norm reconstruction loss is more suitable for retaining the original properties in feature elimination task.

C. MELANOMA CLASSIFICATION

1) DATA

The ISIC 2020 Challenge Dataset has a severely imbalanced distribution of benign and malignant data where only 1.7630% of the training dataset is malignant. Our constructed dataset with hair labels also has of only 1.7259% of malignant samples. Therefore, the key factor in solving the classification in this case comes down to handling the data imbalance problem. However, the objective of this paper is not to show how to maximize the performance of classification, but to propose the new unsupervised hair removal algorithm and clarify the impact of the proposed hair removal algorithm on the classification. For that purpose, we conducted experiments for classification with two different dataset configurations. First, we test classification performance on the entire constructed dataset with hair labels where the melanoma labels are significantly imbalanced. Additionally, we evaluate classification performance on a portion of the constructed dataset with hair labels, where the portion consists of the same number of samples of malignant and benign data. The entire dataset is divided into 6,396 training, 800 validation and 800 testing splits and the partial dataset is composed of 228 training, 18 validation and 30 testing splits. For the experiments using the classification models of ResNet-18 [62] and EfficientNet-B3 [64], we resized the input resolution to 300×300 and for the cases where we used the other models, we used the input whose resolution is resized to 380×380.

TABLE 2. The average classification metric scores of the top 5 results on test set of label-balanced partial dataset with different classification models.

input	metric	Res18	Res152	RX101	EC-B3	EC-B4
original	acc	0.7400	0.7667	0.7133	0.6467	0.6267
	auc	0.8186	0.7840	0.7556	0.6844	0.6951
	f1	0.7447	**0.7683**	0.7200	0.6503	0.6317
	pre	0.7400	**0.7667**	0.7133	0.6467	0.6267
	rec	0.7387	**0.7663**	0.7110	0.6447	0.6229
	spec	0.7467	0.8000	**0.8000**	0.6533	0.5733
CycleGAN	acc	0.8200	0.7333	0.7267	0.6600	0.7267
	auc	0.8871	0.7884	0.7938	0.7191	0.7538
	f1	0.8231	0.7344	0.7285	0.6651	0.7369
	pre	0.8200	0.7333	0.7267	0.6600	0.7267
	rec	0.8195	0.7329	0.7261	0.6590	0.7234
	spec	**0.8667**	0.7333	0.7067	0.7333	0.6800
UNet+L2	acc	0.7800	0.7600	0.7867	0.7467	0.7267
	auc	0.8107	0.8178	0.8284	0.7716	0.7840
	f1	0.7849	0.7618	**0.7886**	0.7500	0.7296
	pre	0.7800	0.7600	**0.7867**	0.7467	0.7267
	rec	0.7791	0.7596	**0.7863**	0.7459	0.7257
	spec	0.7200	0.8000	**0.8000**	**0.7600**	0.6800
ours	acc	**0.8267**	**0.7800**	**0.8133**	**0.7533**	**0.7444**
	auc	**0.9093**	**0.8196**	**0.8489**	**0.8124**	**0.8111**
	f1	**0.8288**	0.7422	0.7327	**0.7618**	**0.7489**
	pre	**0.8267**	0.7333	0.7267	**0.7533**	**0.7444**
	rec	**0.8263**	0.7305	0.7246	**0.7517**	**0.7433**
	spec	0.8000	**0.8267**	0.7733	0.7067	**0.7111**

TABLE 3. The average classification metric scores of the top 5 results only on hair data of test set using trained models with label-balanced partial dataset.

input	metric	Res18	Res152	RX101	EC-B3	EC-B4
original	acc	0.8222	0.8667	0.8222	0.6667	0.6889
	auc	0.9250	0.9000	0.8750	0.7750	0.8750
	f1	0.6875	0.7333	0.7000	**0.6367**	0.6367
	pre	0.8125	0.9250	0.9000	**0.8125**	0.8250
	rec	0.6946	0.7757	0.7271	0.5903	0.6053
	spec	0.7467	0.8000	**0.8000**	0.6533	0.5733
CycleGAN	acc	0.8444	0.8222	0.7333	0.6667	0.7556
	auc	0.6250	0.8500	0.9000	0.6500	0.9000
	f1	0.4416	0.7000	0.6500	0.5190	**0.7233**
	pre	0.4750	0.9000	0.8500	0.5500	**0.8625**
	rec	0.4574	0.7271	0.6410	0.4943	**0.6957**
	spec	0.8667	0.7333	0.7067	**0.7333**	0.6800
UNet+L2	acc	**0.8667**	0.9111	0.8667	0.7556	0.8000
	auc	**0.9750**	**0.9750**	0.9500	0.7250	0.8750
	f1	**0.7333**	0.7875	**0.7333**	0.5899	0.5917
	pre	**0.9250**	0.8625	**0.9250**	0.6000	0.7125
	rec	**0.7757**	0.8075	**0.7757**	0.5746	0.6064
	spec	0.7200	0.8000	**0.8000**	0.7067	0.6800
ours	acc	0.8444	**0.9333**	**0.8889**	**0.7778**	**0.8148**
	auc	0.9250	**0.9750**	**0.9750**	**0.8500**	**0.9167**
	f1	0.7167	**0.8000**	**0.7333**	0.6292	0.6687
	pre	0.9125	**0.9500**	**0.9250**	0.7875	0.8229
	rec	0.7514	**0.8400**	**0.7757**	**0.6359**	0.6929
	spec	**0.8000**	**0.8267**	0.7733	0.7067	**0.7111**

2) EVALUATION

We compare the classification performance on the original images and the reconstructed images where hair is removed by the benchmarks and our proposed method. For the classification, we use a variety of modern classification deep neural networks including ResNet-18 [62], ResNet-152 [62], ResNeXt-101 [63], EfficientNet-B3 [64] and EfficientNet-B4 [64]. In all of the experiments of classification, we used mini-batch sizes of 8 and trained the melanoma classifier C with a learning rate of 1e-05 for 24,000 iterations. For the best model selection, we evaluated accuracy on validation data every 24 iterations and selected the model with the highest validation accuracy. For the experiments where the reconstructed hair-removed images are used as inputs, if an image is originally hairless, the original image y is used as an

input and if the image initially contains hair x, the image is pre-processed by the pre-trained hair removal reconstructor R and the reconstructed hair-removed image γ is used as the input for the melanoma classification network C as described in Sec. III-C. The results of the entire dataset is shown in Tab. 1 where the all column shows the classification accuracy of the all data and the hair column represents the classification accuracy only on data with hair. The results indicate that eliminating hair features from dermoscopic images increases classification accuracy and our proposed algorithm is most effective among the hair removal models. However, as explained in Sec. IV-C1, the distribution of the melanoma labels is extremely imbalanced in the entire dataset. To clarify the effectiveness of our proposed method, we additionally conducted experiments for classification on the partial dataset where the labels of benign and malignant are equally distributed. We conducted 10 experiments for each classification model with each input, selected the 5 results with the highest test accuracy, and computed the average scores of the results. We evaluate the performance on the label-balanced partial dataset with 5 different classification metrics: accuracy, AUC, F1 score, precision, recall and specificity. The metric scores on the partial dataset are given in Tab. 2 and Tab. 2 where the metric column, acc, auc, f1, pre, rec, and spec denote accuracy, AUC, F1 score, precision, recall and specificity; the original row presents results on the original images where hair removal process is not employed; the CycleGAN, UNet+L2 and ours rows present results on the images where hair is removed by CycleGAN, UNet+L2 and ours; and Res, RX and EC stand for ResNet [62], ResNeXt [63] and EfficientNet [64], respectively. The results show that using an image where hair is removed by our method provides the best scores in most cases as shown in Tab. 2. The results on images with hair demonstrate that removing hair considerably improves the classification performance on hair-containing images and the proposed algorithm is most effective in removing hair features for melanoma classification on images with hair as shown in c. The results verify that applying the proposed hair elimination algorithm significantly enhances the performance of the melanoma classification, outperforming the benchmarks.

V. CONCLUSION

Prior focus on undesirable feature elimination algorithms in computer vision was mainly on denoising, dehazing and deraining. Classical approaches characterize statistical properties of undesirable features and succeeding supervised data-driven methods require fine-grained labels to be used as ground-truth. Recent unsupervised techniques for denoising, dehazing and deraining rely on CycleGAN. Meanwhile, hair is an undesirable feature that causes serious performance degradation in dermoscopy applications but studies on hair removal algorithms employing modern deep learning techniques have not been sufficiently provided. In these regards, we have proposed a new unsupervised algorithm for the removal of undesirable features and demonstrated it by the application to dermoscopic images of melanoma where hair results in detrimental impacts. The proposed method characterizes hair simply by a binary classifier, a discriminator of GAN, in a coarse-grained level and imposes the distribution of hairless images to the reconstructed images via GAN. At the same time, the important features of the original input remain in the reconstructed images by minimizing the L_1-norm reconstruction loss function that is built upon the assumption that the additive noise of the original images follows Laplace distribution. We have qualitatively demonstrated that the proposed algorithm is robust to variations of density, thickness, length of hair and colors of skin and lesion. The quantitative evaluations on melanoma classification show that removing hair using the proposed algorithm significantly improves the classification performance and outperforms the benchmarks, indicating its effectiveness and applicability to diagnosis based on dermoscopy.

REFERENCES

[1] M. Nasir, M. A. Khan, M. Sharif, I. U. Lali, T. Saba, and T. Iqbal, "An improved strategy for skin lesion detection and classification using uniform segmentation and feature selection based approach," *Microsc. Res. Technique*, vol. 81, no. 6, pp. 528–543, Jun. 2018.
[2] M. A. Khan, T. Akram, M. Sharif, A. Shahzad, K. Aurangzeb, M. Alhussein, S. I. Haider, and A. Altamrah, "An implementation of normal distribution based segmentation and entropy controlled features selection for skin lesion detection and classification," *BMC Cancer*, vol. 18, no. 1, pp. 1–20, Dec. 2018.
[3] T. Saba, M. A. Khan, A. Rehman, and S. L. Marie-Sainte, "Region extraction and classification of skin cancer: A heterogeneous framework of deep CNN features fusion and reduction," *J. Med. Syst.*, vol. 43, no. 9, pp. 1–19, Sep. 2019.
[4] S. Gu, L. Zhang, W. Zuo, and X. Feng, "Weighted nuclear norm minimization with application to image denoising," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 2862–2869.
[5] W. Dong, L. Zhang, G. Shi, and X. Li, "Nonlocally centralized sparse representation for image restoration," *IEEE Trans. Image Process.*, vol. 22, no. 4, pp. 1620–1630, Apr. 2013.
[6] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Sep. 2009, pp. 2272–2279.
[7] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 2, Jun. 2005, pp. 60–65.
[8] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
[9] X. Liu, M. Tanaka, and M. Okutomi, "Noise level estimation using weak textured patches of a single noisy image," in *Proc. 19th IEEE Int. Conf. Image Process.*, Sep. 2012, pp. 665–668.
[10] D. Zoran and Y. Weiss, "Scale invariance and noise in natural images," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Sep. 2009, pp. 2209–2216.
[11] C. Liu, W. T. Freeman, R. Szeliski, and S. B. Kang, "Noise estimation from a single image," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 1, Jun. 2006, pp. 901–908.
[12] L.-W. Kang, C.-W. Lin, and Y.-H. Fu, "Automatic single-image-based rain streaks removal via image decomposition," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1742–1755, Apr. 2012.
[13] Y. Luo, Y. Xu, and H. Ji, "Removing rain from a single image via discriminative sparse coding," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 3397–3405.
[14] Y. Li, R. T. Tan, X. Guo, J. Lu, and M. S. Brown, "Rain streak removal using layer priors," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2736–2744.
[15] L. Zhu, C.-W. Fu, D. Lischinski, and P.-A. Heng, "Joint bi-layer optimization for single-image rain streak removal," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2526–2534.

[16] L.-J. Deng, T.-Z. Huang, X.-L. Zhao, and T.-X. Jiang, "A directional global sparse model for single image rain removal," *Appl. Math. Model.*, vol. 59, pp. 662–679, Jul. 2018.
[17] Y. Wang, S. Liu, C. Chen, and B. Zeng, "A hierarchical approach for rain or snow removing in a single color image," *IEEE Trans. Image Process.*, vol. 26, no. 8, pp. 3936–3950, Aug. 2017.
[18] R. Fattal, "Single image dehazing," *ACM Trans. Graph.*, vol. 27, no. 3, pp. 1–9, 2008.
[19] R. T. Tan, "Visibility in bad weather from a single image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
[20] K. He, J. Sun, and X. Tang, "Single image haze removal using dark channel prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 12, pp. 2341–2353, Dec. 2011.
[21] B. Xie, F. Guo, and Z. Cai, "Improved single image dehazing using dark channel prior and multi-scale retinex," in *Proc. Int. Conf. Intell. Syst. Design Eng. Appl.*, vol. 1, Oct. 2010, pp. 848–851.
[22] H. Xu, J. Guo, Q. Liu, and L. Ye, "Fast image dehazing using improved dark channel prior," in *Proc. IEEE Int. Conf. Inf. Sci. Technol.*, Mar. 2012, pp. 663–667.
[23] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
[24] K. Zhang, W. Zuo, and L. Zhang, "FFDNet: Toward a fast and flexible solution for CNN-based image denoising," *IEEE Trans. Image Process.*, vol. 27, no. 9, pp. 4608–4622, Sep. 2018.
[25] W. Yang, R. T. Tan, J. Feng, J. Liu, Z. Guo, and S. Yan, "Deep joint rain detection and removal from a single image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1357–1366.
[26] D. Eigen, D. Krishnan, and R. Fergus, "Restoring an image taken through a window covered with dirt or rain," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 633–640.
[27] X. Fu, J. Huang, X. Ding, Y. Liao, and J. Paisley, "Clearing the skies: A deep network architecture for single-image rain removal," *IEEE Trans. Image Process.*, vol. 26, no. 6, pp. 2944–2956, Jun. 2017.
[28] X. Fu, J. Huang, D. Zeng, Y. Huang, X. Ding, and J. Paisley, "Removing rain from single images via a deep detail network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3855–3863.
[29] G. Li, X. He, W. Zhang, H. Chang, L. Dong, and L. Lin, "Non-locally enhanced encoder-decoder network for single image de-raining," in *Proc. 26th ACM Int. Conf. Multimedia*, Oct. 2018, pp. 1056–1064.
[30] D. Chen, M. He, Q. Fan, J. Liao, L. Zhang, D. Hou, L. Yuan, and G. Hua, "Gated context aggregation network for image dehazing and deraining," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2019, pp. 1375–1383.
[31] X. Liu, Y. Ma, Z. Shi, and J. Chen, "GridDehazeNet: Attention-based multi-scale network for image dehazing," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 7314–7323.
[32] B. Cai, X. Xu, K. Jia, C. Qing, and D. Tao, "DehazeNet: An end-to-end system for single image haze removal," *IEEE Trans. Image Process.*, vol. 25, no. 11, pp. 5187–5198, Nov. 2016.
[33] B. Li, X. Peng, Z. Wang, J. Xu, and D. Feng, "AOD-Net: All-in-one dehazing network," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4770–4778.
[34] W. Ren, S. Liu, H. Zhang, J. Pan, X. Cao, and M.-H. Yang, "Single image dehazing via multi-scale convolutional neural networks," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2016, pp. 154–169.
[35] W. Ren, L. Ma, J. Zhang, J. Pan, X. Cao, W. Liu, and M.-H. Yang, "Gated fusion network for single image dehazing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3253–3261.
[36] H. Zhang, V. Sindagi, and V. M. Patel, "Image de-raining using a conditional generative adversarial network," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 11, pp. 3943–3956, Nov. 2020.
[37] J. Chen, J. Chen, H. Chao, and M. Yang, "Image blind denoising with generative adversarial network based noise modeling," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3155–3164.
[38] L. D. Tran, S. M. Nguyen, and M. Arai, "GAN-based noise model for denoising real images," in *Proc. Asian Conf. Comput. Vis.*, 2020, pp. 1–13.
[39] M. Sharma, A. Verma, and L. Vig, "Learning to clean: A GAN perspective," in *Proc. Asian Conf. Comput. Vis.* Cham, Switzerland: Springer, 2018, pp. 174–185.
[40] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra, Y. Zhang, L. Sun, and G. Wang, "Low-dose CT image denoising using a generative adversarial network with Wasserstein distance and perceptual loss," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1348–1357, Jun. 2018.
[41] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
[42] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2223–2232.
[43] H. S. Park, J. Baek, S. K. You, J. K. Choi, and J. K. Seo, "Unpaired image denoising using a generative adversarial network in X-ray CT," *IEEE Access*, vol. 7, pp. 110414–110425, 2019.
[44] Y. Wei, Z. Zhang, Y. Wang, J. Fan, S. Yan, and M. Wang, "DerainCycleGAN: A simple unsupervised network for single image deraining and rainmaking," 2019, *arXiv:1912.07015*. [Online]. Available: http://arxiv.org/abs/1912.07015
[45] L.-Y. Huang, J.-L. Yin, B.-H. Chen, and S.-Z. Ye, "Towards unsupervised single image dehazing with deep learning," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2019, pp. 2741–2745.
[46] Y. Jin, G. Gao, Q. Liu, and Y. Wang, "Unsupervised conditional disentangle network for image dehazing," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2020, pp. 963–967.
[47] T. Lee, V. Ng, R. Gallagher, A. Coldman, and D. McLean, "Dullrazor: A software approach to hair removal from images," *Comput. Biol. Med.*, vol. 27, no. 6, pp. 533–543, Nov. 1997.
[48] F.-Y. Xie, S.-Y. Qin, Z.-G. Jiang, and R.-S. Meng, "PDE-based unsupervised repair of hair-occluded information in dermoscopy images of melanoma," *Computerized Med. Imag. Graph.*, vol. 33, no. 4, pp. 275–282, Jun. 2009.
[49] K. Korotkov and R. Garcia, "Computerized analysis of pigmented skin lesions: A review," *Artif. Intell. Med.*, vol. 56, no. 2, pp. 69–90, Oct. 2012.
[50] A. Huang, S.-Y. Kwan, W.-Y. Chang, M.-Y. Liu, M.-H. Chi, and G.-S. Chen, "A robust hair segmentation and removal approach for clinical images of skin lesions," in *Proc. 35th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2013, pp. 3315–3318.
[51] M. T. B. Toossi, H. R. Pourreza, H. Zare, M.-H. Sigari, P. Layegh, and A. Azimi, "An effective hair removal algorithm for dermoscopy images," *Skin Res. Technol.*, vol. 19, no. 3, pp. 230–235, Aug. 2013.
[52] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, "Unsupervised pixel-level domain adaptation with generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3722–3731.
[53] P. Bibiloni, M. González-Hidalgo, and S. Massanet, "Skin hair removal in dermoscopic images using soft color morphology," in *Proc. Conf. Artif. Intell. Med. Eur.* Cham, Switzerland: Springer, 2017, pp. 322–326.
[54] L. Talavera-Martínez, P. Bibiloni, and M. González-Hidalgo, "An encoder-decoder CNN for hair removal in dermoscopic images," 2020, *arXiv:2010.05013*. [Online]. Available: http://arxiv.org/abs/2010.05013
[55] S. Ding, J. Zheng, Z. Liu, Y. Zheng, Y. Chen, X. Xu, J. Lu, and J. Xie, "High-resolution dermoscopy image synthesis with conditional generative adversarial networks," *Biomed. Signal Process. Control*, vol. 64, Feb. 2021, Art. no. 102224.
[56] H. Rashid, M. A. Tanveer, and H. A. Khan, "Skin lesion classification using GAN based data augmentation," in *Proc. 41st Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2019, pp. 916–919.
[57] Z. Qin, Z. Liu, P. Zhu, and Y. Xue, "A GAN-based image synthesis method for skin lesion classification," *Comput. Methods Programs Biomed.*, vol. 195, Oct. 2020, Art. no. 105568.
[58] F. Pollastri, F. Bolelli, R. Paredes Palacios, and C. Grana, "Improving skin lesion segmentation with generative adversarial networks," in *Proc. IEEE 31st Int. Symp. Comput.-Based Med. Syst. (CBMS)*, Jun. 2018, pp. 442–443.
[59] F. Pollastri, F. Bolelli, R. Paredes, and C. Grana, "Augmenting data with GANs to segment melanoma skin lesions," *Multimedia Tools Appl.*, vol. 79, nos. 21–22, pp. 15575–15592, Jun. 2020.
[60] Y. Peng, N. Wang, Y. Wang, and M. Wang, "Segmentation of dermoscopy image using adversarial networks," *Multimedia Tools Appl.*, vol. 78, no. 8, pp. 10965–10981, Apr. 2019.
[61] W. Tu, X. Liu, W. Hu, Z. Pan, X. Xu, and B. Li, "Segmentation of lesion in dermoscopy images using dense-residual network with adversarial learning," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2019, pp. 1430–1434.
[62] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.

[63] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1492–1500.
[64] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," 2019, *arXiv:1905.11946*. [Online]. Available: http://arxiv.org/abs/1905.11946
[65] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 7167–7176.
[66] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "Cycada: Cycle-consistent adversarial domain adaptation," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1989–1998.
[67] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4681–4690.
[68] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, Dec. 2018.
[69] C. Han, H. Hayashi, L. Rundo, R. Araki, W. Shimoda, S. Muramatsu, Y. Furukawa, G. Mauri, and H. Nakayama, "GAN-based synthetic brain MR image generation," in *Proc. IEEE 15th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2018, pp. 734–738.
[70] T. Iqbal and H. Ali, "Generative adversarial network for medical images (MI-GAN)," *J. Med. Syst.*, vol. 42, no. 11, p. 231, Nov. 2018.
[71] J. Dong, Y. Cong, G. Sun, Y. Yang, X. Xu, and Z. Ding, "Weakly-supervised cross-domain adaptation for endoscopic lesions segmentation," *IEEE Trans. Circuits Syst. Video Technol.*, early access, Aug. 12, 2020, doi: 10.1109/TCSVT.2020.3016058.
[72] C. Chen, Q. Dou, H. Chen, J. Qin, and P.-A. Heng, "Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 865–872.
[73] X. Dong, Y. Lei, T. Wang, M. Thomas, L. Tang, W. JCurran, T. Liu, and X. Yang, "Automatic multiorgan segmentation in thorax CT images using U-Net-GAN," *Med. Phys.*, vol. 46, no. 5, pp. 2157–2168, 2019.
[74] Q. Yue, X. Luo, Q. Ye, L. Xu, and X. Zhuang, "Cardiac segmentation from LGE MRI using deep neural network incorporating shape and spatial priors," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2019, pp. 559–567.
[75] K. Roth, A. Lucchi, S. Nowozin, and T. Hofmann, "Stabilizing training of generative adversarial networks through regularization," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 2018–2028.
[76] V. Rotemberg *et al.*, "A patient-centric dataset of images and metadata for identifying melanomas using clinical context," 2020, *arXiv:2008.07360*. [Online]. Available: http://arxiv.org/abs/2008.07360
[77] X. Jin, Z. Chen, J. Lin, Z. Chen, and W. Zhou, "Unsupervised single image deraining with self-supervised constraints," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2019, pp. 2761–2765.

DAHYE KIM received the B.B.A. degree in computer science and engineering from Chung-Ang University, South Korea, in 2019, where she is currently pursuing the M.S. degree in computer vision. Her current research interests include computer vision and machine learning.

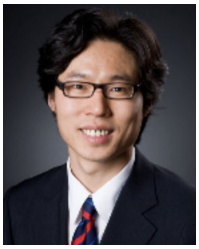

BYUNG-WOO HONG received the M.S. degree in computer vision from the Weizmann Institute of Science with Prof. Shimon Ullman, in 2001, and the Ph.D. degree in computer vision from the University of Oxford with Prof. Michael Brady, in 2005.

In 2008, he joined the Computer Science Department, Chung-Ang University, South Korea, as a Faculty Member, after his postdoctoral research at the Computer Science Department, University of California Los Angeles, with Prof. Stefano Soatto. His research interests include image processing, computer vision, machine learning, and medical image analysis.