Revisiting the Shape-Bias of Deep Learning for Dermoscopic Skin Lesion Classification

Adriano Lucieri[1,2][0000−0003−1473−4745], Fabian Schmeisser[1][0000−0001−8222−7900], Christoph Peter Balada[2][0000−0003−0307−7866], Shoaib Ahmed Siddiqui[2][0000−0003−4600−7331], Andreas Dengel[1,2][0000−0002−6100−8255], and Sheraz Ahmed[2][0000−0002−4239−6520]

[1] Department of Computer Science,
Technische Universität Kaiserslautern,
Erwin-Schrödinger-Straße 52, 67663 Kaiserslautern, Germany
[2] Smart Data and Knowledge Services (SDS),
German Research Center for Artificial Intelligence GmbH (DFKI),
Trippstadter Straße 122, 67663 Kaiserslautern, Germany
`firstname.lastname@dfki.de`

Abstract. It is generally believed that the human visual system is biased towards the recognition of shapes rather than textures. This assumption has led to a growing body of work aiming to align deep models' decision-making processes with the fundamental properties of human vision. The reliance on shape features is primarily expected to improve the robustness of these models under covariate shift. In this paper, we revisit the significance of *shape-biases* for the classification of skin lesion images. Our analysis shows that different skin lesion datasets exhibit varying biases towards individual image features. Interestingly, despite deep feature extractors being inclined towards learning entangled features for skin lesion classification, individual features can still be decoded from this entangled representation. This indicates that these features are still represented in the learnt embedding spaces of the models, but not used for classification. In addition, the spectral analysis of different datasets shows that in contrast to common visual recognition, dermoscopic skin lesion classification, by nature, is reliant on complex feature combinations beyond *shape-bias*. As a natural consequence, shifting away from the prevalent desire of shape-biasing models can even improve skin lesion classifiers in some cases.

Keywords: Dermatology · Digital Dermatoscopy · Skin Lesion Analysis · Spectral Analysis · Robustness · Deep Learning

1 Introduction

For over a decade now, Deep Neural Networks (DNNs) outperformed conventional techniques in various research areas including language translation [26], image classification [17] and image synthesis [11]. Although new state-of-the-art

performances are being reported continuously for areas like Melanoma detection [28,12], an unconstrained application of Deep Learning (DL) in real-world, high-stakes medical decision-making is still considered questionable due to a lack of robustness and intelligibility. Several works have revealed weak spots of the current technology like the presence of adversarial examples [27], the influence of distribution shifts [22] and the bias-variance tradeoff [10].

Geirhos et al. [9] disproved the widespread *shape hypothesis*, which states that Convolutional Neural Networks (CNNs) hierarchically combine lower-lever features into higher-level features for generating the final predictions. Instead the authors propose the *texture hypothesis*, stating that an inherent *texture-bias* in the dataset can lead to a lack of robustness in CNNs. Similarly, other works [32] have reported a higher importance of texture-like high-frequency input features, which aligns with the vulnerability to high-frequency adversarial attacks [30]. Through expensive modification of the training dataset, exchanging the dataset's *texture-bias* to a *shape-bias*, the authors of [9] achieve improved classification robustness. Along similar lines, recent works [5,31] exploit the idea that the phase spectrum of a Fourier-transformed image mainly encodes semantic information resembling edges and outlines used by humans for object identification. The authors propose different data augmentation strategies for improved robustness, inducing explicit focus on the phase spectra of images, shifting the networks' focus towards shape information. The common idea in these works is the explicit alignment of a network's non-functional requirements with those used in analytical, human decision-making (i.e. focusing on shape more than on texture).

Despite recent efforts towards suppressing high-frequency texture features in DL, Ilyas et al. [13] argue that datasets can contain robust features which are indeed imperceptible to humans. This alternative perspective is particularly interesting when dealing with complex medical problems which are yet to be fully understood by human experts and cannot be easily solved through intuition. One of such high-stakes use-cases of DL in medicine is the classification of Melanoma, which is mainly driven by non-analytic clinical reasoning (i.e. pattern analysis [16]).

The statistical relevance of shapes, textures and colors in dermoscopic images for Melanoma detection has been extensively investigated in different studies. Marques et al. [19] reported that color and texture features individually have a high relevance for skin lesion classification, but their combination is even more informative. In other works [24,2] the superior role of color features is reported. Ruela et al. [23] investigate the importance of shape features, concluding that, although shape is relevant for classification, the use of texture and color descriptors is more effective. Beyond texture, shape, and color, other studies indicate a high relevance of spectral features for predictive performance [3,18]. However, the influence of individual features, as well as the effect of the *shape-bias* on DL-based skin lesion classifiers has not yet been explored.

In this paper, we revisit the *shape-bias* and it's effect on the analysis of dermoscopic images using DL-based models. To that end, we explore the relevance

of individual image features known to be relevant in dermoscopy (i.e. *Texture*, *Shape*, and *Color*). A spectral analysis on different datasets is performed to investigate the distribution of relevant image features in the spectral domain, and to revisit the effectiveness of robustness methods enforcing explicit *shape-bias* on deep feature extractors. Lastly, a new variant of the Amplitude-Phase Recombination [5] robustness method is introduced, which is more aligned with the complex needs of dermoscopic skin lesion analysis. We argue that the current trend of focusing network robustness in Deep Learning purely on the *shape-bias* is to narrow-minded, and that medical imaging tasks (like dermoscopy) in particular, have radically different requirements when it comes to non-functional properties of their decision-making.

Section 2 gives a brief introduction into the notions of *Texture*, *Shape*, and *Color* in dermoscopy and describes image ablations used to isolate these different features. The datasets used in throughout our work, as well as the general experimentation setting is outlined in section 3. In section 4, the individual importance of isolated image features is investigated, and their encoding in the DL-based models' feature space. A spectral analysis on different datasets is performed in section 5, followed by an investigation of shape-focused robustness methods. Finally, the results are discussed in section 6, followed by the concluding remarks.

2 Definition and Isolation of Texture, Shape, and Color in Skin Lesions

To properly investigate the influence of individual image features, we need to isolate image features and feature combinations from the input images. We follow the previous lines of work and concentrate on the *Texture*, *Shape* and *Color* features as the main components descriptive of skin lesion images. First, the individual features are briefly defined, based on the relevant literature. Then, the transformations achieving the different feature isolations are elaborated and presented.

Texture Marques et al. [19] define textures in skin lesions as conveying "*information about the differential structures (pigment network, dots, streaks, etc) present in the lesion*". We therefore argue that textures are solely encoded in structures such as fine edges, and color contrasts.

Shape Shape descriptors are computed in [23] based on the segmented lesion outline. The measures include simple shape descriptors such as the lesion's area, compactness, and rectangularity, but also more advanced features such as symmetry-related features and moment invariants. In this work, we define the shape of a lesion by the size and area of a lesion's segmentation, as well as the regularity and overall shape of it's outline. For the sake of simplicity, we omit information regarding the smoothness of a lesion's transition.
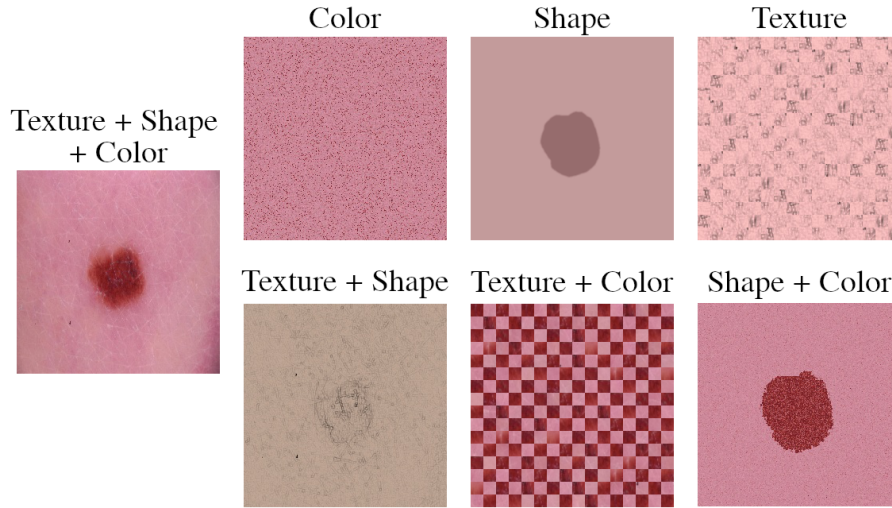
Fig. 1: Illustration of the different augmentations designed to isolate individual image features in skin lesion images. The original image is provided on the left. The first row shows augmentations for isolated images, while the second row shows combinations of two individual features.

Color In [19], color features are defined as containing "*information about the color distribution and number of colors in the skin lesion*". Ruela et al. [24] compute different color descriptors based on different color spaces. Most descriptors contain only information about the quantitative distribution of color in an image, whereas one descriptor also encodes information about the spatial location. We argue that color information is not only encoded in the absolute color values, but also in the contrast information of broader surfaces. However, we do not regard spatial color information for the sake of simplicity.

2.1   Feature Ablations

To isolate the effect of individual features, we design different data augmentation strategies, each representing one of the seven unique feature combinations. An illustration of the different transformations is provided in figure 1. The original images combine all three features and serve as the baseline. An ablation representing only the *Color* feature is obtained by randomly scrambling the spatial ordering of individual image pixels, while preserving the sample's original color distribution. A combination of *Texture* and *Shape* is achieved by explicitly removing the *Color* cues. This is done by changing the style of a sample to a sketch-like image, in order to remove color value and contrast information, while maintaining the characteristic edges necessary to identify textures and shapes.

A DL-based segmentation model[3] is tuned for the computation of lesion segmentation maps, representing the isolated *Shape* feature. The *Shape* feature is removed from the original image with the help of the segmentations as well. Therefore, the segmentation map is divided in equally-sized patches to identify image regions containing information about the lesion's outline (i.e. containing both lesion and background pixels). After removing all outline patches, a new image is assembled by alternatingly sampling random patches from the lesion and the background, to ensure that no information about the lesion's absolute size is retained. A similar procedure is followed to obtain the isolation of the *Texture* feature. Instead of scrambling image patches of the original image, the sketch transformation is used in order to spare color information. A combination of *Shape* and *Color* is obtained by separately scrambling the spatial ordering of individual image pixels within the lesion region, and the background. Images with radical domain shifts (through sketch transformation or segmentation) are shaded with the channel-wise average of the dataset's color.

3   Datasets and Methodology

3.1   Datasets

ISIC & ISIC-b The International Skin Imaging Collaboration (ISIC) organized several skin cancer classification challenges over the last decade. The challenge datasets are hosted on the ISIC's online archive[4], which is the largest public database of dermoscopic skin images to date. The complete archive consists of over 69.000 clinical and dermoscopic images of different provenance. We follow Cassidy et al. [4], who propose a duplicate removal strategy for the ISIC challenge datasets to avoid overlap between training and evaluation sets. For experimentation, we combine all duplicate-free challenge training sets and generate new training, validation and testing splits under stratification.

The complete ISIC dataset has annotations for eight classes, i.e. Actinic Keratosis (AK), Basal Cell Carcinoma (BCC), Benign Keratosis (BKL), Dermatofibroma (DF), Melanoma (MEL), Nevus (NV), Squamous Cell Carcinoma (SCC) and Vascular Lesions (VASC). We generate a multi-class variant of the dataset (henceforth referred to as *ISIC*) comprising of 23.868 training, 2.653 validation and 2.947 testing samples. In addition, a binary variant consisting of only NV and MEL samples (henceforth referred to as *ISIC-b*) is generated comprising of 10.543 training, 4.519 validation and 6.456 testing samples.

D7P & D7P-b The seven-point checklist criteria dataset (*D7P*) proposed in [14] consists of clinical and dermoscopic images of 1.011 skin lesions. Each image is annotated with regards to its diagnostic class, several dermoscopic criteria as well as further clinical data. In this work, we only consider the subset of

---

[3]BA-Transfomer architecture proposed by Wang et al. [29], trained on ISIC2016-2018 challenge datasets.

[4]https://www.isic-archive.com/

dermoscopic images along with the respective annotations of dermoscopic criteria and final diagnosis. We follow the original work, categorizing the fine-grained annotations into BCC, NV, MEL, Sebbhoreic Keratosis (SK) as well as a miscellaneous (MISC) classes. Again, we generate a stratified binary variant of only NV and MEL samples (henceforth referred to as *D7P-b*) comprising of 371 training, 183 validation and 273 testing samples. For concept detection experiments, we follow the same splitting procedure for each dermoscopic concept.

3.2   Experimental Setup

If not mentioned otherwise, all experiments are conducted with a ResNet50, pre-trained on *ImageNet*. Training is conducted using softmax cross-entropy loss and AdamW optimizer. The learning rate and weight decay are determined by hyperparameter tuning on the baseline setting. A plateau learning rate scheduler is used in conjunction with an early stopping scheme to ensure convergence of the models. Each training and respective evaluation is run 10 times with varying random seeds to ensure significance of the reported results.[5]

4   Deep Feature Extractors for Dermoscopy Can Encode Disentangled Features

In this section, we perform an extensive study on the influence of individual features for the DL-based classification of dermoscopic skin lesions. We show that even within the dermoscopy domain, biases are dataset-dependent. Moreover, we show that although feature extractors are inclined towards learning entangled features, last-layer retraining can recover at least some features successfully.

4.1   Different Dermoscopic Skin Lesion Datasets have Different Biases

In a first experiment, we train separate classifiers on the individual ablations introduced in section 2. By training and testing on a specific ablation, we want to quantify the CNN-based feature extractor's capability to leverage individual, interpretable features from the images. We consider both binary and multiclass skin lesion datasets to account for possible effects of varying task complexity.

Figure 2 shows the macro averaged F1-scores on the test set of the respective ablations. When providing only single features in isolation, we observe that both binary and multiclass datasets based on *D7P*, show a stronger bias towards *Shape*, whereas the *ISIC* datasets are more sensitive with respect to *Texture*. This is indicated by the lower decrease in F1-scores when training on the respective isolated feature. It can also be observed that *Color* is the most important of all three features, resulting in the lowest performance decline.

---

[5]Reproducible code available on GitHub https://github.com/adriano-lucieri/shape-bias-in-dermoscopy
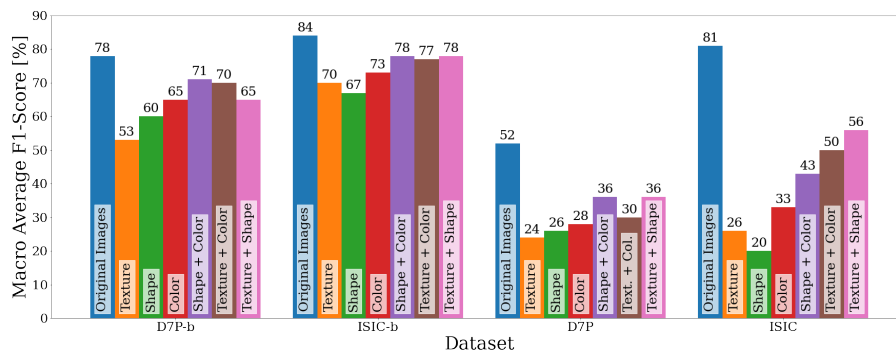
Fig. 2: Macro averaged F1-scores from models trained and tested on individual *Texture*, *Shape* and *Color* feature ablations. The first bar of each group represents the reference F1-score achieved by models trained on unaltered data, followed by the three feature isolation and feature removal ablations.

Training on a pair of two individual features can be considered as the removal of the absent feature, and therefore serves as an inverted indicator for feature importance. The previous observations are also confirmed by the results of removing *Texture* and *Shape*, except for *ISIC-b* which is influenced almost equally. Surprisingly, removing the *Color* did not indicate similar relevance as indicated by the isolation experiment. A reason for this behaviour could be that the combination of *Shape* and *Texture* information forms stronger higher-level features as compared to combinations including *Color*.

4.2   Dermoscopic Skin Lesion Classifiers Learn Entangled Features

The previous results showed that for some datasets, decent classification performances can be achieved even if only one or two features are present in the data. We now investigate the ability of trained classifiers to transfer their features in ablated scenarios. Therefore, the performance of each classifier trained in the previous experiment is measured across all ablations, as well as the original data.

Figure 3 shows a comparison of the results obtained when evaluating the models trained on the original input data on the test sets of all individual ablations (blue bars), with the inverse case, where different models trained on individual ablations are evaluated only on the original test set (orange bars). The data clearly shows that the baseline classifier is unable to properly transfer it's learned features to the classification of ablated data, representing isolated input features. However, the increased scores of the ablation-trained classifiers on the original data indicates the validity of the features even on unablated data. Hence, we conjecture that the baseline models are not capable of making decisions purely based on the remaining features, but are instead overrelying on an entangled representation of different features.
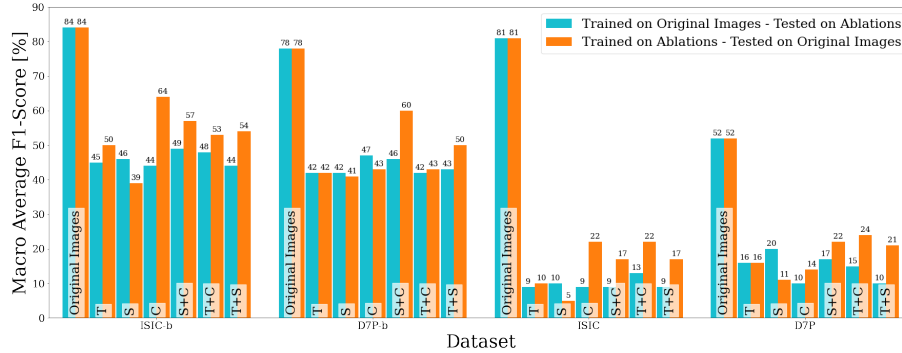
Fig. 3: Comparison of macro averaged F1-scores from models trained and tested on different data ablations. Blue bars show results of models trained on original images, and tested on different ablations. Orange bars show results of models trained on individual ablations, and tested on original images. *T*, *S*, *C* refer to *Texture*, *Shape* and *Color* features, respectively.
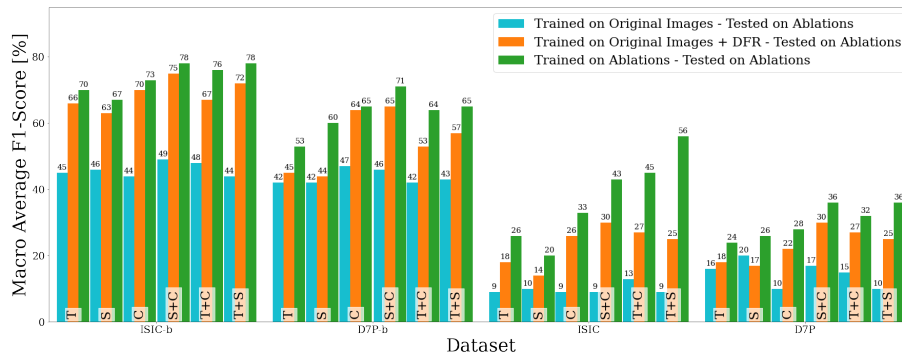

Fig. 4: Comparison of macro averaged F1-scores from models tested on different feature ablations. Blue bars show results of a baseline model trained on original images, and tested on all ablations. Orange bars show results of the same models after *Deep Feature Reweighting*. Green bars show results of models trained and tested on ablations. *T*, *S*, *C* refer to *Texture*, *Shape* and *Color* features, respectively.

4.3   Feature Extractors are only partially Feature Biased

Inspired by *Deep Feature Reweighting* (DFR) proposed in [15], we explore the level of entanglement in the skin lesion classifiers' feature spaces. We again utilise the models trained on the original images from section 4.1 but retrain the fully connected classification layers using the respective feature-ablated training and test sets. In figure 4, the macro averaged F1-scores of baseline models with naive

transfer, and DFR is compared to the accuracies of the models trained end-to-end on the ablated data.

It can be seen that in some cases (e.g. *Color*, *Shape + Color* and partially *Texture + Color* as well as *Texture + Shape* for *ISIC-b*) DFR is able to achieve classification performances near the ideal values represented by the results of models trained and tested on ablations. For models trained on different skin lesion datasets, the successfully recovered features vary significantly. In contrast to *D7P* and *D7P-b* models, where only *Color* could be recovered, models trained on *ISIC-b* were able to recover all individual features to a sufficient degree. A similar observation can be made when inspecting the results of multiclass *ISIC* trained models. Another striking observation is that combinations of two features resulted in higher DFR performance across all datasets. The results indicate that the feature extractor is always inclined towards learning entangled features. However, the fact that individual features are additionally encoded, particularly in more complex and feature-rich datasets, confirms the results reported in [15]. This also suggests that an abundance of mostly redundant features in a dataset allows networks to learn alternative, isolated representations.

5   Dermoscopy Relies on Complex Feature Combinations in Spectral Domain

In this section, we investigate the difference between feature entanglement in skin lesion classification tasks and common visual recognition. For comparability, we consider the distinction of features in the spectral domain, which has been commonly examined in previous studies [32,5,31]. We show that, compared to conventional visual recognition tasks, dermoscopy is more reliant on features across both amplitude and phase spectra. To contrast the spread of dermoscopic features to those used in common visual recognition tasks, we utilize two subsets of *ImageNet* [6], namely *Imagenette* and *Imagewoof* [6].

5.1   Dermoscopy Features are Spread over Phase and Amplitude

First, we implement phase- and amplitude-randomization augmentations which are applied to train, validation and test images alike. Phase-randomization (i.e. *Amplitude-Only*) is applied by replacing the phase spectrum of an image, with the phase spectrum of a randomly sampled image of Gaussian noise in Fourier domain. The same procedure is followed analogously for Amplitude-randomization (i.e. *Phase-Only*).

Table 1 shows the test results of models trained in the *Amplitude-* and *Phase-Only* settings in comparison to the baseline classification accuracies over five different datasets. It can be observed that *Amplitude-Only* always leads to comparatively high deterioration of accuracy for all datasets. However, *Phase-Only* results in comparable accuracy drops over all skin lesion datasets, even causing a higher relative performance decrease for *D7P-b*. Both visual recognition

---

[6]https://github.com/fastai/imagenette

Table 1: Resulting testing accuracies from retraining ResNet50 with different spectral randomizations. For both spectral randomizations, the performance decrease with respect to the baseline training is provided in a separate row.

Dataset	D7P-b	D7P	ISIC-b	ISIC	Imagenette	Imagewoof
Baseline	82.01	69.13	88.45	85.93	97.77	92.53
Amplitude-Only	69.71	36.53	78.42	52.94	50.28	26.86
Δ Baseline	-15.00	-47.16	-11.34	-38.39	-48.57	-70.97
Phase-Only	67.66	45.45	80.78	69.29	91.98	81.53
Δ Baseline	-17.50	-34.25	-08.67	-19.36	-05.92	-11.89

datasets instead show a significantly higher decrease in accuracy when providing only amplitude information, as compared to the *Phase-Only* setting.

The results indicate that skin lesion datasets rely heavily on both amplitude- and phase-spectra, therefore potentially considering a more complex composition and variety of features beyond simple shape information. In contrast, both *ImageNet* subsets show a significant bias towards the phase spectra of the images. The higher drop in accuracy when training *Imagewoof Phase-Only* indicates that a combination of phase and color is extremely important to achieve high performance in some classes, although phase being mostly sufficient. Interestingly, an inspection of the confusion matrices reveals that in the *Phase-Only* setting, networks tend to most often confuse *Beagles* with *English Foxhounds*, which share many features in their physique. A similar observation was made for the *Amplitude-Only* setting, where *Rhodesian Ridgebacks* and *Dingos* where confused most often, sharing a similar fur color.

5.2   Focusing on Amplitude can Improve Performance

Amplitude-Phase Recombination (*APR*) has been proposed in [5] as a method to increase the robustness of DL classifiers by focusing the feature extraction on the phase-spectrum of images. $\mathcal{F}_x$ is the spatial Discrete Fourier Transform (DFT) of an image x over each individual channel. This frequency representation can be decomposed into an amplitude component ($\mathcal{A}_x$) and a phase component ($\mathcal{P}_x$) as follows:

$$\mathcal{F}_x = \mathcal{A}_x \otimes e^{i \cdot \mathcal{P}_x}, \tag{1}$$

Amplitude-Phase Recombination for pair samples (*APR-P*) [5] augments a given input image x_j by replacing its amplitude spectrum $\mathcal{A}_{x_j}$ with the spectrum of another randomly selected image from the batch ($\mathcal{A}_{x_k}$).

$$x_{j,aug} = iDFT(\mathcal{A}_{x_k} \otimes e^{i \cdot \mathcal{P}_{x_j}}), \tag{2}$$

Instead, we propose two variations of *APR-P*, namely Amplitude-Focused *APR-P* (*AF-APR-P*) and Mixed *APR-P* (*Mix-APR-P*). *AF-APR-P* swaps the phase spectrum of images as follows:

$$x_{j,aug} = iDFT(\mathcal{A}_{x_j} \otimes e^{i \cdot \mathcal{P}_{x_k}}), \tag{3}$$

Table 2: Average test results of random retraining with different variants of *APR* augmentation. Statistical significance of average accuracy to the baseline training is indicated by asterisks.

Dataset	D7P-b	ISIC-b	D7P	ISIC	Imagenette	Imagewoof
Baseline	82.01	88.45	**69.13**	**85.93**	**97.77**	**92.53**
APR-P	80.4*	88.62	66.05*	82.79	96.52	91.47
AF-APR-P	**82.42**	**89.23****	67.43	84.18	97.06**	91.91***
MIX-APR-P	81.5	89.16	68.62	84.22	97.3**	91.68

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$

Therefore, the ground truth label corresponding to the original image's amplitude spectrum is preserved. *Mix-APR-P* randomly selects the spectral component from which to assign the respective label, forcing the network to extract both phase and amplitude features.

Table 2 shows the test performance of models trained with different variations of *APR-P* augmentation. *AF-APR-P* outperformed the other variations in the case of binary skin lesion classification, although statistical significance is only achieved in case of *ISIC-b*. When classifying skin lesions in multiple disease classes, neither augmentation showed any benefit. However, it can be seen that *APR-P* decreased the average test accuracy in all skin classification tasks. For both visual recognition datasets, neither *APR* augmentation improved the results. As expected, *AF-APR-P* and *Mix-APR-P* even led to a significant decrease in most cases.

6 Discussion

We have shown there exist different, dataset-dependend biases with respect to *Texture*, *Shape* and *Color* features in dermoscopy. *D7P* datasets seem to be more biased towards *Shape* as compared to *Texture*. One possible explanation for this finding is the significant difference in training size between *ISIC* and *D7P* ($\times 28$ for binary and $\times 53$ for multiclass) allows the *ISIC*-trained models to pick up more nuanced features including fine-grained textures. This effect requires further investigation and is of special importance for the robustness and explainability of skin lesion classifiers in clinical use.

The results in section 4.3 show differences between the feature entanglement of binary and multiclass classification tasks and indicated that a partial or even full encoding of disentangled features is possible depending on the complexity of the target task. However, it has also been shown that the end-to-end trained classifier does not necessarily use these features independently, potentially suffering a loss in robustness. Additional experimentation is required to investigate potential mechanisms leading to disentangled classification layers from end-to-end training. One possible way would be the explicit data augmentation with the different feature isolations proposed in section 2. Another interesting direction

would be the application of contrastive losses from self-supervised learning to achieve a better alignment of different feature-isolated ablations.

Section 5.1 revealed that features relevant for dermoscopic classification are spread across different components in the spectral domain, whereas visual recognition classifiers are deliberately biased towards only the phase component in order to increase robustness. A reason for this phenomenon might lie in the relevance of *Texture*, *Shape* and *Color* in skin lesion classification. The phase spectrum is known to encode mainly edges, which often correspond to coarse structures as outlines, but also fine-grained structures resembling textures. On the other hand, color is mainly encoded in the amplitude spectrum, as it contains information about the magnitude of specific frequency components in the respective color channels.

These findings indicated that common methods for increased robustness, which reinforce the *shape-bias* are not necessarily suitable for skin lesion classification. However, skin lesion classification has been shown to significantly benefit from our proposed Amplitude-Focused *APR* augmentation. Multiclass skin lesion classification seem to not benefit, or even suffer from the Fourier-domain augmentation. A reason for this behaviour might be the increasing relevance of *Texture* and *Shape* features in multiclass settings, as reported in section 4.3. The intuition behind the *Mix-APR-P* augmentation was that a random exposure to phase- or amplitude-randomized samples might implicitly force the network to learn individual features. Yet, this assumption has been disproved.

Overall, the findings of this work indicate an inherent complexity of the dermoscopic skin lesion classification task. Indeed, the process of clinical reasoning has already been shown to be fundamentally different from other human decision such as visual recognition, based mostly on analytical reasoning. Norman et al. [21] describe the process of clinical reasoning as an iterative approach, combining non-analytical with analytical operations to varying degrees, depending upon personal style preferences, experience, and awareness of the diagnostic task [7]. Commonly applied diagnostic procedures in dermoscopy are manual algorithms like the ABCD rule [25], the Menzies method [20], the seven-point checklist [1] as well as the method of pattern analysis [16]. Methods like the seven-point checklist and pattern analysis are based on the identification of complex dermoscopic features such as Blue-Whitish-Veil or Pigment Networks. The successful application of pattern analysis requires years of extensive training. Evidence suggest, that clinical reasoning in dermoscopy puts more emphasis on the non-analytical, unconscious description of overall patterns as compared to analytical processes [8,33].

This suggests that dermoscopic skin lesion classification and other medical imaging tasks pose particularly interesting challenges upon the whole computer vision community. Due to the inherent complexity of the task, special efforts are required from the community working on explainable Artificial Intelligence, in order to properly disentangle features and align explanations with the human clinical reasoning processes.

7   Conclusion

In this paper, we revisit the utility of developing *shape-biased* models for recognition beyond natural images. Particularly, we consider the domain of dermoscopic skin lesion classification. Through a range of different experiments, we have shown that deep features learnt for the classification of skin lesions are inherently entangled due to the complexity of the underlying task. At the same time, our analysis reveals that feature disentanglement can be achieved even on networks trained without constraints, and found that an increasing task complexity as well as a higher number of training samples leads deep feature extractors to learn a more diverse set of redundant and isolated features. Additionally, we showed that dermoscopic features are spread over different spectral components in contrast to common visual recognition tasks like *Imagenet*. This indicates that the commonly desired *shape-bias* for improved model robustness does not apply in dermoscopy, and that the task requires specifically tailored solutions. We demonstrated a first step towards dermoscopy-specific robustness measures beyond *shape-bias* by introducing Amplitude-Focused Amplitude-Phase Recombination, showing improved performance on binary skin lesion classification tasks. More importantly, this work highlights the importance of scrutinizing a given computer vision task in order to find relevant, and robust requirements for the decision-making. Dermoscopy is only one out of plenty use-cases with unique requirements which extend beyond the simple analytical procedures of visual object recognition. These kind of considerations are particularly important for pivotal areas such as self-supervised learning, where an adequate requirement engineering will potentially lead to enormous performance improvements.

References

1. Argenziano, G., Fabbrocini, G., Carli, P., De Giorgi, V., Sammarco, E., Delfino, M.: Epiluminescence microscopy for the diagnosis of doubtful melanocytic skin lesions: comparison of the abcd rule of dermatoscopy and a new 7-point checklist based on pattern analysis. Archives of dermatology **134**(12), 1563–1570 (1998)
2. Barata, C., Ruela, M., Mendonça, T., Marques, J.S.: A bag-of-features approach for the classification of melanomas in dermoscopy images: The role of color and texture descriptors. In: Computer vision techniques for the diagnosis of skin cancer, pp. 49–69. Springer (2014)
3. Betta, G., Di Leo, G., Fabbrocini, G., Paolillo, A., Sommella, P.: Dermoscopic image-analysis system: estimation of atypical pigment network and atypical vascular pattern. In: IEEE International Workshop on Medical Measurement and Applications, 2006. MeMea 2006. pp. 63–67. IEEE (2006)
4. Cassidy, B., Kendrick, C., Brodzicki, A., Jaworek-Korjakowska, J., Yap, M.H.: Analysis of the isic image datasets: usage, benchmarks and recommendations. Medical image analysis **75**, 102305 (2022)
5. Chen, G., Peng, P., Ma, L., Li, J., Du, L., Tian, Y.: Amplitude-phase recombination: Rethinking robustness of convolutional neural networks in frequency domain. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 458–467 (2021)

6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)

7. Dinnes, J., Deeks, J.J., Chuchu, N., Matin, R.N., Wong, K.Y., Aldridge, R.B., Durack, A., Gulati, A., Chan, S.A., Johnston, L., et al.: Visual inspection and dermoscopy, alone or in combination, for diagnosing keratinocyte skin cancers in adults. Cochrane Database of Systematic Reviews (12) (2018)

8. Gachon, J., Beaulieu, P., Sei, J.F., Gouvernet, J., Claudel, J.P., Lemaitre, M., Richard, M.A., Grob, J.J.: First prospective study of the recognition process of melanoma in dermatological practice. Archives of dermatology **141**(4), 434–438 (2005)

9. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. arXiv preprint arXiv:1811.12231 (2018)

10. Geman, S., Bienenstock, E., Doursat, R.: Neural networks and the bias/variance dilemma. Neural computation **4**(1), 1–58 (1992)

11. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems **27** (2014)

12. Hasan, M.K., Elahi, M.T.E., Alam, M.A., Jawad, M.T., Martí, R.: Dermoexpert: Skin lesion classification using a hybrid convolutional neural network through segmentation, transfer learning, and augmentation. Informatics in Medicine Unlocked p. 100819 (2022)

13. Ilyas, A., Santurkar, S., Tsipras, D., Engstrom, L., Tran, B., Madry, A.: Adversarial examples are not bugs, they are features. Advances in neural information processing systems **32** (2019)

14. Kawahara, J., Daneshvar, S., Argenziano, G., Hamarneh, G.: Seven-point checklist and skin lesion classification using multitask multimodal neural nets. IEEE Journal of Biomedical and Health Informatics **23**(2), 538–546 (mar 2019). https://doi.org/10.1109/JBHI.2018.2824327

15. Kirichenko, P., Izmailov, P., Wilson, A.G.: Last layer re-training is sufficient for robustness to spurious correlations. arXiv preprint arXiv:2204.02937 (2022)

16. Kittler, H., Rosendahl, C., Cameron, A., Tschandl, P.: Dermatoscopy: an algorithmic method based on pattern analysis. facultas (2016)

17. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems **25** (2012)

18. López-Leyva, J.A., Guerra-Rosas, E., Álvarez-Borrego, J.: Multi-class diagnosis of skin lesions using the fourier spectral information of images on additive color model by artificial neural network. IEEE Access **9**, 35207–35216 (2021)

19. Marques, J.S., Barata, C., Mendonça, T.: On the role of texture and color in the classification of dermoscopy images. In: 2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. pp. 4402–4405. IEEE (2012)

20. Menzies, S.W., Ingvar, C., Crotty, K.A., McCarthy, W.H.: Frequency and morphologic characteristics of invasive melanomas lacking specific surface microscopic features. Archives of dermatology **132**(10), 1178–1182 (1996)

21. Norman, G., Barraclough, K., Dolovich, L., Price, D.: Iterative diagnosis. Bmj **339** (2009)

22. Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D.: Dataset shift in machine learning. Mit Press (2008)

23. Ruela, M., Barata, C., Mendonca, T., Marques, J.S.: On the role of shape in the detection of melanomas. In: 2013 8th International Symposium on Image and Signal Processing and Analysis (ISPA). pp. 268–273. IEEE (2013)
24. Ruela, M., Barata, C., Mendonça, T., Marques, J.S.: What is the role of color in dermoscopy analysis? In: Iberian Conference on Pattern Recognition and Image Analysis. pp. 819–826. Springer (2013)
25. Stolz, W.: Abcd rule of dermatoscopy: a new practical method for early recognition of malignant melanoma. Eur. J. Dermatol. **4**, 521–527 (1994)
26. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. Advances in neural information processing systems **27** (2014)
27. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. In: International Conference on Learning Representations (2014), http://arxiv.org/abs/1312.6199
28. Tang, P., Yan, X., Nan, Y., Xiang, S., Krammer, S., Lasser, T.: Fusionm4net: A multi-stage multi-modal learning algorithm for multi-label skin lesion classification. Medical Image Analysis **76**, 102307 (2022)
29. Wang, J., Wei, L., Wang, L., Zhou, Q., Zhu, L., Qin, J.: Boundary-aware transformers for skin lesion segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 206–216. Springer (2021)
30. Wang, Z., Yang, Y., Shrivastava, A., Rawal, V., Ding, Z.: Towards frequency-based explanation for robust cnn. arXiv preprint arXiv:2005.03141 (2020)
31. Xu, Q., Zhang, R., Zhang, Y., Wang, Y., Tian, Q.: A fourier-based framework for domain generalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14383–14392 (2021)
32. Yin, D., Gontijo Lopes, R., Shlens, J., Cubuk, E.D., Gilmer, J.: A fourier perspective on model robustness in computer vision. Advances in Neural Information Processing Systems **32** (2019)
33. Zalaudek, I., Kittler, H., Marghoob, A.A., Balato, A., Blum, A., Dalle, S., Ferrara, G., Fink-Puches, R., Giorgio, C.M., Hofmann-Wellenhof, R., et al.: Time required for a complete skin examination with and without dermoscopy: a prospective, randomized multicenter study. Archives of dermatology **144**(4), 509–513 (2008)