

Secondary Publication



Di Salvo, Francesco; Doerrich, Sebastian; Ledig, Christian

MedMNIST-C : Comprehensive benchmark and improved classifier robustness by simulating realistic image corruptions

Date of secondary publication: 06.08.2025

Submitted Version (Preprint), Conferenceobject

Persistent identifier: urn:nbn:de:bvb:473-irb-109532x

Primary publication

Di Salvo, Francesco; Doerrich, Sebastian; Ledig, Christian (2024): MedMNIST-C : Comprehensive benchmark and improved classifier robustness by simulating realistic image corruptions, ADSMI : MICCAI Workshop on Advancing Data Solutions in Medical Imaging AI, October 10, 2024 ; Marrakesh, Morocco, pp. 1–11, doi: 10.48550/arXiv.2406.17536.

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by/4.0/legalcode>

MedMNIST-C: Comprehensive benchmark and improved classifier robustness by simulating realistic image corruptions

Francesco Di Salvo, Sebastian Doerrich, and Christian Ledig

xAILab Bamberg, University of Bamberg, Germany
francesco.di-salvo@uni-bamberg.de

Abstract. The integration of neural-network-based systems into clinical practice is limited by challenges related to domain generalization and robustness. The computer vision community established benchmarks such as ImageNet-C as a fundamental prerequisite to measure progress towards those challenges. Similar datasets are largely absent in the medical imaging community which lacks a comprehensive benchmark that spans across imaging modalities and applications. To address this gap, we create and open-source **MedMNIST-C**, a benchmark dataset based on the MedMNIST+ collection covering 12 datasets and 9 imaging modalities. We simulate task and modality-specific image corruptions of varying severity to comprehensively evaluate the robustness of established algorithms against real-world artifacts and distribution shifts. We further provide quantitative evidence that our simple-to-use artificial corruptions allow for highly performant, lightweight data augmentation to enhance model robustness. Unlike traditional, generic augmentation strategies, our approach leverages domain knowledge, exhibiting significantly higher robustness when compared to widely adopted methods. By introducing **MedMNIST-C** and open-sourcing the corresponding library allowing for targeted data augmentations, we contribute to the development of increasingly robust methods tailored to the challenges of medical imaging. The code is available at github.com/francescodisalvo05/medmnistc-api.

Keywords: Robustness · Dataset · Benchmark · Data augmentation.

1 Introduction

In the past decade, Deep Neural Networks (DNNs) have made significant improvements, achieving impressive results in various domains, including healthcare. As such the medical imaging community has witnessed a substantial rise of deep learning based decision-support systems, leading to significant advancements in fields such as radiology [28], dermatology [19], and pathology [9], often improving human-performances. However, traditional neural networks continue to face challenges, notably with adversarial samples and distribution shifts [6]. These shifts commonly occur as they are caused by multiple factors including imaging machines (*e.g.*, different vendors), post-processing techniques, patient characteristics, and acquisition protocols.

While the natural imaging domain has advanced in addressing these issues through works such as the ImageNet-C benchmark [10], the medical imaging domain has seen fragmented efforts, with individual works tackling specific modalities, as dermatology [18], digital pathology [33,13], blood microscopy [32], and more. Additionally, they adopt diverse evaluation metrics, such as accuracy [32], average corruption error (CE), corruption error of confidence (CEC) [33], normalized CE, relative normalized CE [13], and their *balanced* counterparts [18]. These works, while foundational, highlight the absence of a unified evaluation setting that can address the diversity of medical imaging modalities.

Unanimously, they confirm the lack of robustness of widely adopted deep learning algorithms. To address model robustness in a simple way, we typically rely on generic augmentation techniques, such as MixUp [31] (blending examples and labels), CutMix [30] (interchanging image patches), RandAugment [4] (chain of random policies), and AugMix [11] (extends RandAugment through multiple chains and a consistency loss). However, those methods are not consistently effective across datasets. A recent work [7] showed superior performance of *targeted augmentation* over generic and domain-invariant ones. While showing promising results on distribution shifts for histopathological images, it lacks a wider validation for the medical domain.

Motivated by the unresolved issues and the potential of targeted augmentation, our work carves out a new path. We first introduce **MedMNIST-C**, a comprehensive image corruption benchmark for 12 datasets and 9 imaging modalities, leveraging the publicly available MedMNIST+ [29] dataset collection (resolution: 224×224). The designed corruptions are dataset-specific and reflect the possible artifacts encountered during image acquisition and processing, simulating real-world artifacts or possible distribution shifts. Then, we employ the designed corruptions as custom *targeted augmentations*. The proposed augmentations can be readily applied via our open-source APIs during training and meaningfully enhance the robustness of traditional domain-agnostic approaches. In summary:

- We *design targeted image corruptions* to simulate real-world artifacts across the 12 datasets and 9 imaging modalities included in the MedMNIST+ dataset collection.
- We *introduce a novel robustness benchmark*, bridging the widely utilized MedMNIST+ dataset with the well-established robustness evaluation framework of ImageNet-C. This provides a comprehensive framework to assess the robustness of algorithms across diverse imaging modalities.
- We *evaluate the robustness of widely adopted deep learning architectures* using our benchmark, highlighting the need for enhancing model robustness.
- We *demonstrate the effectiveness of our targeted augmentations*, increasing the robustness through a simpler and more interpretable method, compared to traditional ones.
- We *publicly share our corrupted datasets and APIs* to facilitate the online, and dataset-specific augmentation of medical image analysis problems.

2 The MedMNIST-C dataset

2.1 Corruptions

This section summarizes the relevant literature and outlines the motivation for our chosen corruptions. As summarized in Table 1, we divided the corruptions into five categories, each with specific corruptions applicable to the respective MedMNIST+ (test) dataset. **Digital** corruptions are applied to any dataset and include *JPEG* compression and *pixelate*, the latter mimicking the effect of upsampling low-resolution images. The other categories are **noise**, **blur**, **color**, and **task-specific**. Following Hendrycks et al. [10] and subsequent medical imaging benchmarks [18,33,13], we evaluate model robustness across 5 *severity levels*. Therefore, we carefully reviewed, identified, and then described how those severity levels can be mapped to a broad set of medical imaging datasets, as illustrated in Figure 1.




Fig. 1: Overview of four different corruptions applied (from top to bottom) to PathMNIST, ChestMNIST, DermaMNIST, and RetinaMNIST.

PathMNIST and BloodMNIST Pathology and blood cell microscopy share similar imaging protocols, leading to common artifacts [33,32,13]. These include *stain deposits* and *air bubbles*, along with *defocus* and *motion blur* during image acquisition. Variations in *brightness*, *contrast*, and *saturation* are also prevalent, arising from different illumination and scanner conditions.

Dataset	Corruption categories				
	Digital	Noise	Blur	Color	Task-specific
PathMNIST	JPEG	-	Defocus	Brightness _[+ -]	Stain deposit
BloodMNIST	Pixelate	-	Motion	Contrast _[+ -] Saturate	Bubble
ChestMNIST		Gaussian			
PneumoniaMNIST	JPEG	Speckle	Gaussian	Brightness _[+ -]	Gamma corr. _[+ -]
OrganAMNIST	Pixelate	Impulse		Contrast _[+ -]	
OrganCMNIST		Shot			
OrganSMNIST					
DermaMNIST	JPEG Pixelate	Gaussian Speckle Impulse Shot	Defocus Motion Zoom	Brightness _[+ -] Contrast _[+ -]	Black corner Characters
RetinaMNIST	JPEG Pixelate	Gaussian Speckle	Defocus Motion	Brightness _[-] Contrast _[-]	-
TissueMNIST	JPEG Pixelate	Impulse	Gaussian	Brightness _[+ -] Contrast _[+ -]	-
OCTMNIST	JPEG Pixelate	Speckle	Defocus Motion	Contrast _[-]	-
BreastMNIST	JPEG Pixelate	Speckle	Motion	Brightness _[+ -] Contrast _[-]	-

Table 1: Overview of the selected image corruptions, each applied at *five* increasing severity levels. Note that the corruption hyperparameters are defined at dataset-level. The symbols $[+|-]$ indicate whether the corruption intensity is increased, decreased, or both, yielding separate corruptions.

PneumoniaMNIST, ChestMNIST, and OrganMNIST (A,C,S) Chest X-Ray and abdominal CT utilize X-ray radiation, with CT scans offering greater detail at higher radiation doses. Key corruptions [22,21,15] include *brightness* and *contrast* variations, and *Gaussian blur*. Additionally, X-ray images are particularly prone to *Gaussian*, *speckle*, *impulse*, and *shot* noise. Lastly, both modalities often undergo *gamma correction* to adjust luminance.

DermaMNIST In dermatoscopy, skin lesions are examined using a dermatoscope. Supported by a comprehensive prior corruption benchmark [18] and further works like [15], we propose: noise artifacts (*Gaussian*, *speckle*, *impulse*, *shot*), blurring effects (*defocus*, *motion*, *zoom*), and color-based artifacts (*brightness*, *contrast*). Finally, we also include task-specific artifacts like *black corners* from the dermatoscope and *characters* from camera overlays.

RetinaMNIST Fundus photography is essential for diagnosing retinal diseases. Here, image quality is affected by optimal *brightness* and *contrast*, influenced by lighting [24]. To reflect diagnostic challenges, we also include *Gaussian* and *speckle* noise for electronic or sensor irregularities [3]. Additionally, blurs like *Gaussian* and *defocus* are also frequent [2].

TissueMNIST High-throughput microscopy, which is characterized by a high image acquisition rate, often experiences various image artifacts. Common issues include *impulse noise* [34] and *blur* (Gaussian) [20]. Additionally, uneven illumination [27] is a frequent problem. This might affect the *brightness* and *contrast* of the images, resulting in inconsistent image quality.

OCTMNIST Optimal Coherence Tomography (OCT) is an imaging modality providing high-resolution images of the retina. OCT images commonly exhibit several artifacts [1,17], including *speckle noise*, *motion blur*, and *defocus blur*. Notably, *speckle noise* arises from the limited spatial-frequency bandwidth of the signal. This noise reduces the *contrast* of the images, thereby degrading their quality and potentially affecting diagnostic accuracy.

BreastMNIST Breast ultrasound uses sound waves to visualize breast tissue and detect abnormalities. Frequently observed artifacts [25,14] consist of *speckle noise*, due to the scattering of ultrasound waves. Additionally, we observe variation in *brightness*, low *contrast*, and *motion blur*.

2.2 Robustness measures

The MedMNIST+ datasets span four distinct tasks: binary classification, multi-class classification, multi-label (binary) classification, and ordinal regression. Following [29], we treat ordinal regression tasks as multi-class classification problems, maintaining consistency across our evaluations. To develop our evaluation setting, we draw inspiration from the ImageNet-C benchmark [10], which employs *corruption error* and *relative corruption error* as key metrics, subsequently adopted in [13] in the context of digital pathology. Those metrics are rooted in the *error* notation, intuitively defined as $1.0 - acc$. Nevertheless, considering the diverse imbalance ratios across our datasets, we follow the approach of [18] for skin cancer, using the *balanced error* (*i.e.*, $1.0 - bacc$).

Thus, we first assess the clean balanced error, BE_{clean}^f , for each model f using the respective MedMNIST+ test set. Next, we evaluate the balanced error, denoted as $BE_{s,c}^f$, for each corruption $c \in C_d$ and severity s (from 1 to 5). Note that C_d indicates the set of all corruptions associated with dataset d (*e.g.*, $C_{blood} = \{JPEG, \dots, brightness+, brightness-, \dots, bubble\}$). Then, we average across severities, normalizing those errors with AlexNet’s errors, yielding BE_c^f . This step, formalized in Equation 1, accounts for the varying impacts of

different corruptions on classification performance. Subsequently, we derive BE^f by averaging across corruptions.

$$\text{BE}_c^f = \frac{\sum_{s=1}^5 \text{BE}_{s,c}^f}{\sum_{s=1}^5 \text{BE}_{s,c}^{\text{AlexNet}}} \quad (1)$$

Moreover, we measure the relative balanced error rBE to assess the performance drop with respect to the clean test set. This is crucial to evaluate model’s robustness, as we usually aim for minimal performance drop under distribution shifts. To do so, we simply extend Equation 1 by subtracting the clean performance.

$$\text{rBE}_c^f = \frac{\sum_{s=1}^5 (\text{BE}_{s,c}^f - \text{BE}_{\text{clean}}^f)}{\sum_{s=1}^5 (\text{BE}_{s,c}^{\text{AlexNet}} - \text{BE}_{\text{clean}}^{\text{AlexNet}})} \quad (2)$$

Finally, we average rBE_c^f across all corruptions to derive rBE^f . For multi-label tasks, like ChestMNIST, metrics are averaged across labels as well.

2.3 Corruption-API for Data Augmentation

The data augmentation method employs the corruptions outlined in Section 2.1. Specifically, for a given dataset d , we uniformly choose *one* random corruption c from the extended set $C'_d = C_d \cup \{\text{identity}\}$, where C_d defines the corruptions associated with the imaging modality of dataset d , and **identity** represents no corruption. Upon selecting a specific corruption, we also randomly determine its severity. Note that each corruption’s severity level is linked to distinct hyperparameters. For instance, with the *brightness+* corruption, the intensity (i) might range from 1.1 at severity level 1 to 1.9 at level 5. Thus, we randomly sample the corruption’s hyperparameter using a uniform distribution that spans from the parameters defined for severity 1 and 5 (*e.g.*, $i \sim U(1.1, 1.9)$).

3 Experimental results

Robustness The first set of experiments aims to demonstrate the effect of our proposed target-specific corruptions on the performance of commonly used deep learning models. Specifically, we evaluate: AlexNet [16] (as normalizing baseline, following [10]), ResNet50 [8], DenseNet121 [12], ViT-B/16 [5], and VGG16 [26]. All networks were initialized with ImageNet weights and trained on the respective MedMNIST+ training sets (*cf.* Table 3 for number of training samples) for 100 epochs with early stopping based on the respective validation set. The AdamW optimizer with a learning rate of $1e^{-4}$ was utilized, with a cosine annealing learning rate schedule. To evaluate the robustness we use the **BE** and the **rBE** described in Section 2.2, averaging across three seed runs.

Table 2 presents a broad overview, averaging the model performance across all the proposed datasets. As expected [23], while ViT-B/16 is the most robust

one, VGG follows in terms of normalized performance (BE). However, despite achieving the best *clean performance*, probably due to the highest number of parameters, it exhibits a severe performance drop (rBE). Notably, ResNet50 is the least robust one, followed by DenseNet121. This goes against the findings of the ImageNet-C benchmark. Thus, the training set size might play an important role. The table also shows that the degree of robustness is different among corruption types. For instance, *task-specific* corruptions have the least impact on ResNet50, DenseNet121, and ViT. Also, while ViT exhibits higher robustness to all corruptions, the gap is even more pronounced on *noise* and *digital* artifacts.

Arch	#Par	M+	M-C		BE ↓					
		bACC ↑	bACC ↑	rBE ↓	BE ↓	Digital	Noise	Blur	Color	TS
AlexNet	62.3 M	78.7	62.9	100.0	100.0	100	100	100	100	100
R.Net50	25.6 M	75.4	56.2	166.1	131.5	177	110	123	148	95
D.Net121	8 M	79.8	59.4	148.4	114.8	145	124	100	124	78
VGG16	138.4 M	80.5	65.9	114.0	93.0	128	87	91	84	80
ViT-B	86.6	78.9	72.0	59.9	76.3	74	50	77	80	71

Table 2: Balanced accuracy, BE, and rBE, averaged across all 12 datasets included in our benchmark, spanning from low-data settings (<1,000 train samples, *e.g.*, BreastMNIST) to large-data ones (>100,000 train samples, *e.g.*, TissueMNIST). Note that #Par refers to the number of parameters, M+ refers to the clean MedMNIST+ test set, and M-C refers to our corrupted test set.

Data augmentation The second experiment aims to evaluate our augmentation method against reference ones, on the corrupted test sets, utilizing a ResNet18 for consistency. We benchmark our method against generic augmentation approaches: MixUp ($\alpha = 0.2$), CutMix ($\alpha = 1.0$), and RandAugment ($k = 1$), with α values chosen from recommended ranges [7]. RandAugment, sharing similarities with our approach through augmentations like *color*, *contrast*, *brightness*, and *sharpness*, differs by employing domain-agnostic intensity ranges. To fairly compare our method, we limit RandAugment to a single corruption at a time ($k = 1$). For the same reason, we exclude from our comparison AugMix, which builds upon RandAugment with multiple chains. In future work, the extension of our targeted augmentations along multiple chains could be explored. While the generic augmentations are employed once and tested on the corrupted test sets, we apply a double evaluation for our method. Initially, considered as a performance baseline, we train and test using a k -fold cross-validation, with folds organized by corruption types: **digital**, **noise**, **blur**, **color**, and **task-specific**. Thus, by training and evaluation with semantically different corruptions, we carefully ensure no overfitting. Secondly, we evaluate using all corruptions, similar to competing methods. Focusing on absolute gains, we report the Area Under the Curve (AUC), averaging across three seed runs.

Table 3 shows that, even under a k -fold cross-validation method, our augmentation approach achieves the top average AUC gain across all 12 datasets. Notably, our augmentation excels particularly in smaller datasets, those with fewer than 10,000 samples, which also tend to have lower initial AUC values. In contrast, without the k -fold settings, we clearly observe significant improvements. While this might seem as an overestimate of performance gains, it pinpoints a sweet spot for our method’s effectiveness, marking substantial enhancements over generic domain-agnostic approaches. This underscores the importance of integrating domain knowledge into data augmentation strategies.

Dataset	Br	Re	Pn	De	Bl	OrC	OrS	OrA	Ch	Pa	OCT	Ti	AVG
#Train	546	1,080	5,856	7,007	11,959	12,975	13,932	34,561	78,468	89,996	97,477	165,466	
#Classes	2	5	2	7	8	11	11	11	14	9	4	8	
No Aug.	73.4	70.8	84.7	78.1	97.8	91.5	89.9	92.0	66.2	92.5	89.3	82.8	84.1
CutMix	1.2	2.4	7.1	-0.7	-0.8	1.1	2.2	0.4	(0)	-1.5	-0.8	0.4	1.0
MixUp	1.6	0.8	5.1	3.0	0.2	5.0	2.9	4.8	(0)	3.4	2.0	1.0	2.7
RandAug.	6.7	3.2	8.7	4.1	1.3	4.0	4.8	2.2	3.2	4.3	3.4	0.8	3.9
Ours (k -F.)	8.5	3.2	7.9	6.8	0.4	2.2	3.9	2.1	4.7	3.2	4.5	0.9	4.0
Ours	12.7	8.5	12.6	16.3	2.1	7.6	7.8	7.4	10.1	7.2	9.2	9.3	9.2

Table 3: AUC gains ($\Delta \uparrow$) achieved by a ResNet18 on the corrupted test sets, employing various augmentation strategies. k -F refers to the k -fold evaluation setting, and the top two strategies are **bolded**. *Ours* is overall the best performing method ($p < 0.05$, Wilcoxon signed-rank test). Note that CutMix and MixUp involve batch-wise transformations and do not natively support multi-label problems. Thus, they do not contribute to the overall mean.

4 Conclusions

This study introduces MedMNIST-C, a comprehensive benchmark specifically designed for assessing algorithm robustness in the context of image corruptions across a wide range of the medical imaging spectrum. We demonstrate the significant advantage of embedding domain knowledge into data augmentation strategies. Our approach, straightforward yet impactful, consistently enhances robustness and outperforms standard methods, even under k -fold validation, ensuring an unbiased evaluation against our test corruptions. There is potential to extend our method to simultaneously manage multiple corruptions, mirroring RandAugment, or to implement them across multiple chains as in AugMix. By providing reproducible corrupted test sets and APIs, we aim to encourage the community to both *assess* and *enhance* the robustness of medical image analysis models.

Acknowledgments. This study was funded through the Hightech Agenda Bayern (HTA) of the Free State of Bavaria, Germany.

References

1. Adabi, S., Fotouhi, A., Xu, Q., Daveluy, S., Mehregan, D., Podoleanu, A., Nasiri-avanaki, M.: An overview of methods to mitigate artifacts in optical coherence tomography imaging of the skin. *Skin Research and Technology* **24**(2) (2018)
2. Araujo, T., Aresta, G., Mendonça, L., Penas, S., Maia, C., Carneiro, Â., Mendonça, A.M., Campilho, A.: Dr| graduate: Uncertainty-aware deep learning-based diabetic retinopathy grading in eye fundus images. *Medical Image Analysis* **63** (2020)
3. Bala, A.A., Priya, P.A., Maik, V.: Retinal image enhancement using curvelet based sigmoid mapping of histogram equalization. In: *Journal of Physics: Conference Series*. vol. 1964, p. 062034. IOP Publishing (2021)
4. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops* (2020)
5. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
6. Drenkow, N., Sani, N., Shpitser, I., Unberath, M.: A systematic review of robustness in deep learning for computer vision: Mind the gap? *arXiv preprint arXiv:2112.00639* (2021)
7. Gao, I., Sagawa, S., Koh, P.W., Hashimoto, T., Liang, P.: Out-of-domain robustness via targeted augmentations. *arXiv preprint arXiv:2302.11861* (2023)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Conference on Computer Vision and Pattern Recognition*. pp. 770–778 (2016)
9. Hekler, A., Utikal, J.S., Enk, A.H., Berking, C., Klode, J., Schadendorf, D., Jansen, P., Franklin, C., Holland-Letz, T., Krahl, D., et al.: Pathologist-level classification of histopathological melanoma images with deep neural networks. *European Journal of Cancer* **115**, 79–83 (2019)
10. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261* (2019)
11. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: AugMix: A simple data processing method to improve robustness and uncertainty. *Proceedings of the International Conference on Learning Representations* (2020)
12. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4700–4708 (2017)
13. Huang, P., Zhang, S., Gan, Y., Xu, R., Zhu, R., Qin, W., Guo, L., Jiang, S., Luo, L.: Assessing and enhancing robustness of deep learning models with corruption emulation in digital pathology. *arXiv preprint arXiv:2310.20427* (2023)
14. Huang, Q., Luo, Y., Zhang, Q.: Breast ultrasound image segmentation: a survey. *International journal of computer assisted radiology and surgery* **12** (2017)
15. Islam, M., Li, Z., Glocker, B.: Robustness stress testing in medical image classification. In: *International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging*. pp. 167–176. Springer (2023)

16. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Neural Information Processing Systems* **25** (2012)
17. Lian, J., Hou, S., Sui, X., Xu, F., Zheng, Y.: Deblurring retinal optical coherence tomography via a convolutional neural network with anisotropic and double convolution layer. *IET Computer Vision* **12**(6), 900–907 (2018)
18. Maron, R.C., Schlager, J.G., Haggemüller, S., von Kalle, C., Utikal, J.S., Meier, F., Gellrich, F.F., Hobelsberger, S., Hauschild, A., French, L., et al.: A benchmark for neural network robustness in skin cancer classification. *European Journal of Cancer* **155**, 191–199 (2021)
19. Maron, R.C., Weichenthal, M., Utikal, J.S., Hekler, A., Berking, C., Hauschild, A., Enk, A.H., Haferkamp, S., Klode, J., Schadendorf, D., et al.: Systematic outperformance of 112 dermatologists in multiclass skin cancer image classification by convolutional neural networks. *European Journal of Cancer* **119**, 57–65 (2019)
20. McElliott, M.C., Al-Suraimi, A., Telang, A.C., Ference-Salo, J.T., Chowdhury, M., Soofi, A., Dressler, G.R., Beamish, J.A.: High-throughput image analysis with deep learning captures heterogeneity and spatial relationships after kidney injury. *Scientific Reports* **13**(1), 6361 (2023)
21. Momeny, M., Neshat, A.A., Hussain, M.A., Kia, S., Marhamati, M., Jahanbakhshi, A., Hamarneh, G.: Learning-to-augment strategy using noisy and denoised data: Improving generalizability of deep cnn for the detection of covid-19 in x-ray images. *Computers in Biology and Medicine* **136**, 104704 (2021)
22. Navarro, F., Watanabe, C., Shit, S., Sekuboyina, A., Peeken, J.C., Combs, S.E., Menze, B.H.: Evaluating the robustness of self-supervised learning in medical imaging. *arXiv preprint arXiv:2105.06986* (2021)
23. Paul, S., Chen, P.Y.: Vision transformers are robust learners. In: *Proceedings of the AAAI conference on Artificial Intelligence*. vol. 36, pp. 2071–2081 (2022)
24. Paulus, J., Meier, J., Bock, R., Hornegger, J., Michelson, G.: Automated quality assessment of retinal fundus photos. *International journal of computer assisted radiology and surgery* **5**, 557–564 (2010)
25. Sehgal, C.M., Weinstein, S.P., Arger, P.H., Conant, E.F.: A review of breast ultrasound. *Journal of mammary gland biology and neoplasia* **11**, 113–123 (2006)
26. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
27. Singh, S., Bray, M.A., Jones, T., Carpenter, A.: Pipeline for illumination correction of images for high-throughput microscopy. *Journal of microscopy* **256**(3) (2014)
28. Wong, S., Al-Hasani, H., Alam, Z., Alam, A.: Artificial intelligence in radiology: how will we be affected? *European radiology* **29**, 141–143 (2019)
29. Yang, J., Shi, R., Ni, B.: Medmnist classification decathlon: A lightweight automl benchmark for medical image analysis. In: *IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. pp. 191–195 (2021)
30. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6023–6032 (2019)
31. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412* (2017)
32. Zhang, S., Ni, Q., Li, B., Jiang, S., Cai, W., Chen, H., Luo, L.: Corruption-robust enhancement of deep neural networks for classification of peripheral blood smear images. In: *Medical Image Computing and Computer Assisted Intervention*. pp. 372–381. Springer (2020)

33. Zhang, Y., Sun, Y., Li, H., Zheng, S., Zhu, C., Yang, L.: Benchmarking the robustness of deep neural networks to common corruptions in digital pathology. In: *Medical Image Computing and Computer-Assisted Intervention (2022)*
34. Zhou, X., Wong, S.: Informatics challenges of high-throughput microscopy. *IEEE Signal Processing Magazine* **23**(3), 63–72 (2006)