



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Shortcut Learning in Medical Image Segmentation

Manxi Lin^{1†}, Nina Weng^{1†}, Kamil Mikolaj¹, Zahra Bashir^{2,3},
Morten B. S. Svendsen^{1,3}, Martin G. Tolsgaard^{3,4,5},
Anders N. Christensen¹, Aasa Feragen¹ (✉)

¹ Technical University of Denmark, Kongens Lyngby, Denmark

{manli, ninwe, afhar}@dtu.dk

² Slagelse Hospital, Slagelse, Denmark

³ CAMES, Copenhagen, Denmark

⁴ Copenhagen University Hospital Rigshospitalet, Copenhagen, Denmark

⁵ Department of Clinical Medicine, University of Copenhagen, Copenhagen, Denmark

Abstract. Shortcut learning is a phenomenon where machine learning models prioritize learning simple, potentially misleading cues from data that do not generalize well beyond the training set. While existing research primarily investigates this in the realm of image classification, this study extends the exploration of shortcut learning into medical image segmentation. We demonstrate that clinical annotations such as calipers, and the combination of zero-padded convolutions and center-cropped training sets in the dataset can inadvertently serve as shortcuts, impacting segmentation accuracy. We identify and evaluate the shortcut learning on two different but common medical image segmentation tasks. In addition, we suggest strategies to mitigate the influence of shortcut learning and improve the generalizability of the segmentation models. By uncovering the presence and implications of shortcuts in medical image segmentation, we provide insights and methodologies for evaluating and overcoming this pervasive challenge and call for attention in the community for shortcuts in segmentation. Our code is public at https://github.com/nina-weng/shortcut_skinseg.

Keywords: Shortcut Learning · Medical Image Segmentation

1 Introduction

Shortcut learning, a critical challenge in machine learning, refers to situations where models latch onto easy-to-learn correlations between input data and target features that are available at training time, but not necessarily at test time. This can result in models that perform well during development, but falter when exposed to real-world practice, where the “shortcuts” are absent. Shortcut learning receives increasing attention in medical imaging [21,23], where it affects high-stakes, high-risk clinical decisions [13], undermines the trust in machine learning models [8], and even threatens model fairness [5].

[†] M. Lin and N. Weng contributed equally to this work.

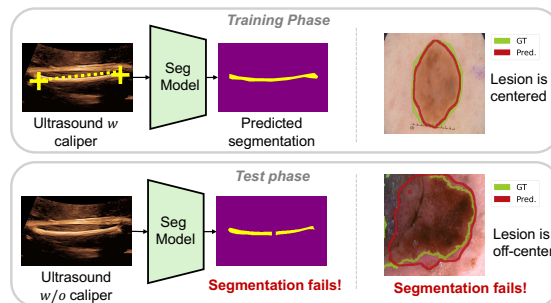


Fig. 1: Left: Shortcuts as yellow markers that affect ultrasound segmentation. The calipers were dilated for illustration. Right: Shortcuts that stem from the dataset construction, where lesions are center-cropped.

Most current work addresses shortcut learning in classification. Fig. 1 shows an example of how clinical ultrasound screening images are, as a natural part of the clinical process, marked with text and measurement calipers at acquisition time. The text and calipers correlate strongly with specific anatomical standard planes, and as demonstrated in [19], this correlation enables the model to use the markers rather than actual anatomical characteristics when classifying images into different anatomical planes. This can cause failure at test time, when e.g. sorting images during an ultrasound examination before calipers are placed.

Shortcut learning has rarely been investigated in other tasks such as image segmentation – an essential medical image quantification task that demands exceptional accuracy and generalizability. Intuitively, shortcut learning might seem less of a risk for image segmentation: In contrast to classification, segmentation necessitates a precise delineation of organ boundaries. This task utilizes spatial relationships and textures within the image, as opposed to standalone, identifiable markers. As a result, it might seem unrealistic that a segmentation model affected by shortcuts could segment anything at all¹. However, this only addresses the danger of false positives.

In this paper, we challenge this prevailing assumption and show that shortcut learning can also affect medical image segmentation models. We show how shortcut learning can lead to false negatives (failure to segment) at test time, greatly affecting segmentation performance.

We showcase two types of shortcuts in segmentation, with different underlying mechanisms, in fetal ultrasound and dermatological segmentation tasks. The first type mirrors the shortcuts identified in classification tasks. The other arises when center-cropped images are zero-padded in CNNs. As a result, the models learn that the objects to be segmented are located in the middle of the image, and they therefore fail to recognize objects near the border. This is important since center cropped patches are a very common segmentation dataset design, also in MICCAI challenges [26,4], see also Fig. 1.

¹ See e.g. <https://gwern.net/tank>, point 6 under “Could it happen”

2 Background and Related Work

A wide range of variables have been shown to act as shortcuts in medical images. In a large collection of examples, the model **learns signs of treatment, diagnosis notes, or signs of related diseases** as proxies of the target y . Known examples include chest tube in chest x-ray as a common treatment for Pneumothorax [13]; clinical annotations or diagnosis notes which encode clinicians’ suspected diagnoses, such as ink circling suspicious skin lesions or rulers measuring their size [21,23]; text and calipers in ultrasound [19] annotated primarily on high-quality images; or effects of disease symptoms such as movement artifact in MRI scans, which can appear as a result of tremor associated with Parkinson’s disease, or decreased capability of following instructions of staying still due to diseases such as dementia [11].

In a different class of shortcuts, the **model encodes global information about the data collection process**, such as a hospital or hardware information, e.g different hospitals might use different scanners and protocols, which give the images a distinctive ‘style’ [10,27]; data collected from different hospitals might have different prevalence distributions, e.g. a cancer-specialized hospital might have more samples of diseased subjects than a community hospital.

Model mistakenly learns the demographic features of the patient, which is associated with the disease prevalence. Studies have shown that models can easily learn the race or gender information from chest X-rays [9], increasing the concerns of using demographic features for prediction, especially when prevalence varies among groups.

Shortcuts performance differs by the choice of model structures. Studies show that model structures and loss functions could impact shortcut learning, e.g., Izmailov and Kirichenko et al. [12] studied how different choices of model architecture and pre-training strategy impact shortcut learning.

Whereas most case studies are classification tasks, in this paper, we showcase that shortcuts can also affect medical image segmentation.

3 Shortcut A: Calipers and texts in fetal ultrasound

Organ segmentation in fetal ultrasound is crucial for measurements and guiding predictive models. Segmentation models are typically trained on clinical screening images, which often include clinical annotations like text and measurement calipers specific to each anatomical plane [14]. These annotations are absent during real-time applications, such as highlighting organs during scans. Here, we document how text and calipers serve as shortcuts for ultrasound segmentation.

3.1 Experiments and observations

Data and experiment settings Our study involved segmenting organs from 3,775 third-trimester ultrasound images derived from a proprietary fetal ultrasound screening dataset. These images represent 13 distinct fetal anatomical

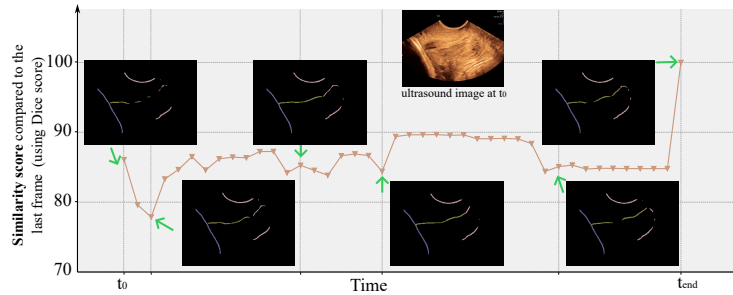


Fig. 2: Illustration of the dynamic variation of model segmentation prediction on the maternal cervix in a video. The similarity score between frames at time t_0 and the last frame t_{end} is computed by dice score. The selected frame shows that the segmentation prediction is unstable when putting the calipers and texts.

structures and originate from four specific anatomical views: the fetal head, femur, abdomen, and the maternal cervix. A medical doctor annotated each image, with an experienced expert conducting subsequent verification. We partitioned the data into training, validation, and test sets, maintaining an 80/10/10 percent split. We utilized the official implementation of DTU-Net [17]. We trained the model with the AdamW optimizer for 1,000 epochs with weight decay $1e-6$, batch size 32, and early stopping. Training began with a learning rate of $1e-4$ and was adjusted by a cosine annealing schedule.

Demonstrating shortcut learning We first validate model performance on two versions of the test set. First, we use original images with text and calipers, and second, a “clean” version of the same test set where annotations are artificially removed via the inpainting strategy from [19]. Tab. 1 shows a performance drop on the clean images, indicating that annotations are used as shortcuts.

We further evaluate the impact of these shortcuts in real-time scans in the clinics. Here, we analyze 110 screen capture video recordings of ultrasound examinations from a single hospital, where the sonographer sequentially identified the anatomical structures of the head, abdomen, and femur in 100 videos, and performed cervix examination in the remaining 10 videos. Upon locating each anatomical feature, the image was frozen and then annotated with texts and

Table 1: Model performance on a test set with clinical annotations, and with the same annotations artificially removed. We report the average and standard deviation of the per-image Dice coefficient.

| Image | Anatomical plane | | | |
|--------------------------------|------------------|------------------|------------------|------------------|
| | Head | Abdomen | Femur | Cervix |
| Test set <i>w</i> annotation | 76.97 ± 5.10 | 82.06 ± 6.60 | 93.82 ± 1.79 | 76.29 ± 4.05 |
| Test set <i>w/o</i> annotation | 70.85 ± 8.24 | 78.85 ± 7.72 | 91.84 ± 4.87 | 71.81 ± 4.86 |

Table 2: Model performance varies during the recorded annotation of a frozen ultrasound plane. The variation is measured with average dice coefficients (%). ‘Baseline’ is the model that suffers from shortcuts while ‘Mitigation’ is the model trained on ‘clean’ images. The mitigated model is better and far more stable.

| Model | Anatomical plane | | | |
|------------|------------------|--------------|--------------|--------------|
| | Head | Abdomen | Femur | Cervix |
| Baseline | 89.07 ± 14.9 | 94.75 ± 12.6 | 91.51 ± 19.6 | 82.42 ± 10.1 |
| Mitigation | 98.39 ± 1.3 | 98.53 ± 1.9 | 95.38 ± 8.9 | 97.00 ± 1.2 |

calipers. We observe fluctuations in the segmentation predictions throughout this process, even when the underlying ultrasound image is frozen, see Fig. 2, and a video clip example can be found in the supplementary material. The first row of Tab. 2 shows the impact of shortcuts on all 110 videos. We quantify segmentation robustness by reporting the difference in average dice coefficients of the segmentation predictions on the frames at the start and the end of the annotation process, i.e., the same image with and without clinical notes.

Mechanism We are inspired by the conclusion from [19] for image recognition. In our segmentation task, the shortcut arises from the consistent co-occurrence of anatomical structures and the calipers and texts within our training dataset. This co-occurrence leads our segmentation model to rely on these clinical notes as predictive shortcuts, rather than fully understanding the texture and shape features of the anatomical structures themselves. This reliance undermines the model’s performance in applications, posing a risk to its clinical applicability.

Mitigation Following [19], we trained a mitigated model on images with annotations artificially removed. As shown in the second row of Tab. 2, performance increases substantially, indicating that text and calipers acted as shortcuts.

4 Shortcut B: The combination of zero-padded convolutions and center cropped training sets

Skin lesions refer to damaged or abnormal skin tissue that differs in appearance from the surrounding area. Common causes of skin lesions include abscesses, cysts, tumors, sores, and ulcers [22]. Accurately segmenting skin lesions is critical for various downstream tasks, such as dermoscopic feature detection and lesion classification, particularly for detecting malignant skin cancers.

Here, we utilized the widely-used ISIC skin lesion segmentation dataset [6] to illustrate our second case of segmentation shortcut learning, which stems from the combination of two typical design choices in image segmentation: (i) In skin lesion segmentation datasets, regions of interests, i.e., lesions tend to be shown quite close to the center of the image. This has been dubbed the ‘common photographer bias’ in computer vision [16] but is also prevalent in



Fig. 3: Sample lesion images from **ISIC_{CROP}** with ground truth (GT) segmentation and predicted segmentation by M_{ori} and M_{crop} . We clearly see shortcut learning and its mitigation when training on M_{ori} and M_{crop} , respectively.




Fig. 4: To emphasize the effect of shortcut learning, we separate the image into yellow bands of increasing distance from the center (a) and compute segmentation Dice scores within each region. Note how the Dice score of the shortcut model M_{ori} declines with increasing distance from the center (b), while the mitigated model M_{crop} maintains its segmentation ability across the image (c).

popular medical datasets, see Fig. 6. **(ii)** A very common modeling choice is to use zero padding in CNN segmentation architectures to create output predictions with the same shape as the input image. In addition to ease of implementation, this also enables a full segmentation, rather than the partial segmentation of the image that a CNN without padding would deliver.

However, the above two design choices combined lead to an unfortunate correlation between the features of pixels near the boundary and their class (typically background): For any pixel whose receptive field extends beyond the original image, its receptive field will contain a band of constant zero values (Fig. 5), and pixels with such features will – due to the centering of lesions in the dataset – almost always belong to the background class. In other words, the combination of zero padding and dataset construction constitutes a shortcut that lets the model learn that pixels near the image boundary should be segmented as background.




Fig. 5: Zero padding introduces a change in the features of segmented pixels whose receptive field contains padding – i.e. those that are near the boundary. This affects their encoding in feature space.

4.1 Experiments and observations

Data and experimental settings We use the **ISIC2017** skin lesion dataset [6] with its original train and test splits. To show the effect of shortcut learning via center cropping and zero padding, we create an asymmetrically cropped version, **ISIC_{CROP}**, where the original images are cropped into 4 quarters, one of which is randomly selected (see Fig. 3). Segmentations were made with a standard UNet [24] trained with Dice loss [25] using Adam [15] with a learning rate of 0.001, 100 epochs and batch size 16. All images and masks are resized to 256×256 .

Demonstrating shortcut learning Segmentation models M_{ori} and M_{crop} were trained using the original train splits within **ISIC2017** and **ISIC_{CROP}**, respectively. As illustrated in Fig 3, M_{ori} has problems segmenting lesions near the boundary. This is further quantified in Fig. 4 (a), where we show how segmentation performance decreases in regions near the image boundary.

Mechanism To understand why zero padding causes shortcut learning, we need to return to the basics of CNNs: CNN segmentation networks are translation equivalent, meaning that without zero padding, the CNN does not encode any information about the spatial location of individual pixels. However, the feature patterns induced by padding enable the CNN to encode – in a band around the image border – the distance to the image border. As shown in Fig. 5, this likely leads to these pixels clustering in feature space due to their special zero padding features. When using the **ISIC2017** dataset to train M_{ori} , those pixels almost exclusively belong to the background, which encourages the model to simply classify all boundary pixels as background.

Mitigation There is a simple yet efficient mitigation strategy: Namely to *not* center crop training images (equivalently, augment with random cropping). This is documented by the M_{crop} model, see Fig. 3 and 4 (c). When lesions are no longer centered during training, test performance no longer depends on the distance to the boundary, and the shortcut has been mitigated. Indeed, while




Fig. 6: Segmentation mask center distributions illustrate center cropping within common medical image segmentation benchmark- and challenge datasets.

M_{ori} can barely categorize any edging pixels into the class of skin lesion, M_{crop} clearly accepts the boundary of the image being the skin lesion.

An alternative mitigation could be to swap zero padding with reflection or replication. As these would still have distinctive features, shortcut learning would still be possible, but likely less pronounced as the shortcuts would be more difficult to learn. Nevertheless, crop augmentation is preferable as it completely removes the correlation between boundary affinity and segmentation labels.

5 Discussion and conclusion

We have shown that shortcut learning is indeed possible for medical image segmentation, expanding the current discourse which focuses narrowly on classification. We demonstrate two different shortcut mechanisms. The ultrasound case mirrors known shortcut mechanisms from image classification, where models use text and calipers to guide segmentation decisions. Inpainting the calipers can mitigate this. Conversely, the interplay of zero padding and center cropped training sets reveals a different type of shortcut, where modeling choices contribute to changing the feature representations of certain pixels. In this context, we note just how prominent center cropping is in popular benchmarks for medical image segmentation (Fig. 6): it seems likely that shortcut learning may be a factor in affecting scores in, for instance, segmentation challenges.

Shortcut learning versus overfitting. Overfitting occurs when excessively complex models capture not just meaningful patterns but also random noise in the training data. This makes the model highly specific to its training set and less capable of generalizing to new data. In contrast, shortcuts lead models to learn superficial and misleading patterns that don't truly represent the underlying problem. The result is still models that perform well on the training data while failing to generalize at test time, but the mechanism is wholly different.

Domain adaptation does not mitigate shortcut learning. Domain adaptation implicitly assumes that the patterns learned by the algorithm on the original domain are transferrable to the target domain. In shortcut learning, however, the algorithm has not necessarily learned the underlying patterns at all, because the shortcuts have provided a simpler strategy for optimizing the training loss. As a result, it might not be possible to apply domain adaptation to transfer a model that has learned shortcuts, to data without shortcuts.

Center-cropped shortcuts are beyond CNNs. Recently popular vision transformers (ViTs) [7] apply no padding and process image patches, which could mitigate shortcuts as the spatial location of the pixel is no longer included. However, the positional encoding in ViTs might still encode the center bias and lead to a similar shortcut, leaving the shortcut in ViTs an interesting open problem.

Beware! Segmentation models are vulnerable to shortcuts. We aim to highlight that shortcut learning risks extend beyond classification, affecting pixel-level tasks like segmentation, detection, super-resolution, denoising, and artifact removal. Addressing these shortcuts is crucial for developing precise, robust, and reliable machine-learning models for clinical use. We show how simple augmentation and inpainting can mitigate our cases of shortcut learning, while less than 15% of skin lesion segmentation papers since 2014 used random cropping augmentation [20], and removing calipers in fetal ultrasound segmentation has not been a standard preprocessing either [3,28]. Future research should address more generic tools to detect and mitigate shortcuts in image segmentation.

Acknowledgments. This work was supported by the DIREC project EXPLAIN-ME (9142-00001B), the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (NNF20OC0062606), and the Pioneer Centre for AI, DNRF grant nr P1. The AI Signature co-funded this work.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Antonelli, M., Reinke, A., Bakas, S., et al.: The medical segmentation decathlon. *Nature Communications* (2022). <https://doi.org/10.1038/s41467-022-30695-9>
2. Armato III, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Zhao, B., Aberle, D.R., Henschke, C.I., Hoffman, E.A., et al.: The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics* **38**(2), 915–931 (2011)
3. Bano, S., Dromey, B., Vasconcelos, F., Napolitano, R., David, A.L., Peebles, D.M., Stoyanov, D.: Autofb: automating fetal biometry estimation from standard ultrasound planes. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VII* 24. pp. 228–238. Springer (2021)
4. Bao, R., Song, Y., Bates, S.V., Weiss, R.J., Foster, A.N., Cobos, C.J., Sotardi, S., Zhang, Y., Gollub, R.L., Grant, P.E., et al.: Boston neonatal brain injury dataset for hypoxic ischemic encephalopathy (bonbid-hie): Part i. mri and manual lesion annotation. *bioRxiv* pp. 2023–06 (2023)
5. Brown, A., Tomasev, N., Freyberg, J., Liu, Y., Karthikesalingam, A., Schrouff, J.: Detecting shortcut learning for fair medical ai using shortcut testing. *Nature Communications* **14**(1), 4314 (2023)
6. Codella, N.C., Gutman, D., Celebi, M.E., Helba, B., Marchetti, M.A., Dusza, S.W., Kalloo, A., Liopyris, K., Mishra, N., Kittler, H., et al.: Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic).

- In: 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018). pp. 168–172. IEEE (2018)
7. 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. In: International Conference on Learning Representations (2020)
 8. Geirhos, R., Jacobsen, J.H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., Wichmann, F.A.: Shortcut learning in deep neural networks. *Nature Machine Intelligence* **2**(11), 665–673 (2020)
 9. Glocker, B., Jones, C., Bernhardt, M., Winzeck, S.: Algorithmic encoding of protected characteristics in chest x-ray disease detection models. *Ebiomedicine* **89** (2023)
 10. Glocker, B., Robinson, R., Castro, D.C., Dou, Q., Konukoglu, E.: Machine learning with multi-site imaging data: An empirical study on the impact of scanner effects. *Medical Imaging Meets NeurIPS Workshop* (2019)
 11. Godenschweger, F., Kägebein, U., Stucht, D., Yarach, U., Sciarra, A., Yakupov, R., Lüsebrink, F., Schulze, P., Speck, O.: Motion correction in mri of the brain. *Physics in medicine & biology* **61**(5), R32 (2016)
 12. Izmailov, P., Kirichenko, P., Gruver, N., Wilson, A.G.: On feature learning in the presence of spurious correlations. *Advances in Neural Information Processing Systems* **35**, 38516–38532 (2022)
 13. Jiménez-Sánchez, A., Juodelyte, D., Chamberlain, B., Cheplygina, V.: Detecting shortcuts in medical images—a case study in chest x-rays. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2023)
 14. Khalil, A., Sotiriadis, A., D’Antonio, F., Da Silva Costa, F., Odibo, A., Prefumo, F., Papageorghiou, A., Salomon, L., et al.: Isuog practice guidelines: performance of third-trimester obstetric ultrasound scan. *Ultrasound in Obstetrics & Gynecology* **63**(1), 131–147 (2024)
 15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
 16. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4015–4026 (2023)
 17. Lin, M., Zepf, K., Christensen, A.N., Bashir, Z., Svendsen, M.B.S., Tolsgaard, M., Feragen, A.: Dtu-net: Learning topological similarity for curvilinear structure segmentation. In: International Conference on Information Processing in Medical Imaging. pp. 654–666. Springer (2023)
 18. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging* **34**(10), 1993–2024 (2014)
 19. Mikolaj, K., Lin, M., Bashir, Z., Svendsen, M.B.S., Tolsgaard, M., Nymark, A., Feragen, A.: Removing confounding information from fetal ultrasound images. arXiv preprint arXiv:2303.13918 (2023)
 20. Mirikharaji, Z., Abhishek, K., Bissoto, A., Barata, C., Avila, S., Valle, E., Celebi, M.E., Hamarneh, G.: A survey on deep learning for skin lesion segmentation. *Medical Image Analysis* p. 102863 (2023)
 21. Narla, A., Kuprel, B., Sarin, K., Novoa, R., Ko, J.: Automated classification of skin lesions: from pixels to practice. *Journal of Investigative Dermatology* **138**(10), 2108–2110 (2018)

22. National Cancer Institute: Lesion. <https://www.cancer.gov/search/results?swKeyword=lesion> (2024), accessed: 2024-03-02
23. Nauta, M., Walsh, R., Dubowski, A., Seifert, C.: Uncovering and correcting shortcut learning in machine learning models for skin cancer diagnosis. *Diagnostics* **12**(1), 40 (2021)
24. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III* 18. pp. 234–241. Springer (2015)
25. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: *DLMIA*. pp. 240–248. Springer (2017)
26. Wang, G., Luo, X., Gu, R., Yang, S., Qu, Y., Zhai, S., Zhao, Q., Li, K., Zhang, S.: Pymic: A deep learning toolkit for annotation-efficient medical image segmentation. *Computer Methods and Programs in Biomedicine* **231**, 107398 (2023)
27. Zech, J.R., Badgeley, M.A., Liu, M., Costa, A.B., Titano, J.J., Oermann, E.K.: Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS medicine* **15**(11), e1002683 (2018)
28. Zhou, S., Ahn, E., Wang, H., Quinton, A., Kennedy, N., Sridar, P., Nanan, R., Kim, J.: Improving automatic fetal biometry measurement with swoosh activation function. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 283–292. Springer (2023)