

SPECTRAL DECOUPLING FOR TRAINING TRANSFERABLE NEURAL NETWORKS IN MEDICAL IMAGING

Joona Pohjonen^{1,*}, Carolin Stürenberg¹, Antti Rannikko^{1,2}, Tuomas Mirtti^{1,3}, and Esa Pitkänen^{4,5,*}

¹Research Program in Systems Oncology, University of Helsinki

²Department of Urology; Helsinki University Hospital

³Department of Pathology; Helsinki University Hospital

⁴Institute for Molecular Medicine Finland (FIMM), HiLIFE

⁵Research Program in Applied Tumor Genomics, University of Helsinki

ABSTRACT

Many current neural networks for medical imaging generalise poorly to data unseen during training. Such behaviour can be caused by networks overfitting easy-to-learn, or statistically dominant, features while disregarding other potentially informative features. For example, indistinguishable differences in the sharpness of the images from two different scanners can degrade the performance of the network significantly. All neural networks intended for clinical practice need to be robust to variation in data caused by differences in imaging equipment, sample preparation and patient populations.

To address these challenges, we evaluate the utility of spectral decoupling as an implicit bias mitigation method. Spectral decoupling encourages the neural network to learn more features by simply regularising the networks' unnormalised prediction scores with an L2 penalty, thus having no added computational costs.

We show that spectral decoupling allows training neural networks on datasets with strong spurious correlations. Networks trained without spectral decoupling do not learn the original task and appear to make false predictions based on the spurious correlations. Spectral decoupling also increases networks' robustness for data distribution shifts. To validate our findings, we train networks with and without spectral decoupling to detect prostate cancer tissue slides and COVID-19 in chest radiographs. Networks trained with spectral decoupling achieve substantially higher performance on all evaluation datasets.

Our results show that spectral decoupling helps with generalisation issues associated with neural networks. We recommend using spectral decoupling as an implicit bias mitigation method in any neural network intended for clinical use.

1 INTRODUCTION

Neural networks have been adapted to many medical imaging tasks with impressive results, often surpassing human counterparts in consistency, speed and accuracy [1]. However, these networks are prone to overfit easy-to-learn, or statistically dominant, features while disregarding other potentially informative features. This leads to poor generalisation to data generated by different medical centres, reliance on the dominant features, and lack of robustness [2, 3]. For example, a neural network classifier for skin cancer, approved to be used as a medical device in Europe, had overfit the correlation between surgical margins and malignant melanoma [4]. Due to this, the false positive rate of the network was increased by 40 percentage points during external validation. Furthermore, three out of five neural networks for pneumonia detection showed significantly worse performance during external validation [5] and recent neural networks for COVID-19 detection rely on confounding factors rather than actual medical pathology [6]. Even small differences in the sharpness of images from two different scanners can degrade the performance of neural networks significantly (see Section 3.2).

Although generalisation issues need to be solved before any neural networks can be applied in clinical practice, the issue is still poorly understood [7]. This may be because the detection of generalisation issues is hard and often requires state-of-the-art methods of explainable AI [6]. Simply using an external validation set will not uncover all generalisation issues as a neural network may achieve high overall accuracy while always failing for some subset of samples, and the external dataset may contain the same sources of bias as the training data.

Some explicit methods have been proposed to address more obvious sources of bias, like using augmentation to address staining differences in tissue section slides [8] or normalising each image with a common standard [9, 10]. The obvious problem with explicit methods is that they only control for selected biases and more subtle sources of bias, like small differences between patient populations, may go unaddressed. Implicit methods of bias control are required before neural networks can be safely applied to clinical practice.

Learning dominant features at the cost of other potentially informative features, also known as shortcut-learning, is a common problem in all neural networks and one of the main reasons behind the generalisation issues [3]. Shortcut-learning occurs mainly because of gradient starvation, where gradient descent updates the parameters of a neural network in directions cap-

*Correspondence:

joona.pohjonen@helsinki.fi

esa.pitkanen@helsinki.fi

turing only dominant features, thus starving the gradient from other features [11]. The gradient descent algorithm finds a local optimum by taking small steps towards the opposite sign of the derivative, the direction of the steepest descent [12]. The recently proposed method of spectral decoupling [2] provably decouples the learning dynamics leading to gradient starvation when using cross-entropy loss, thus encouraging the network to learn more features. The effect is achieved by simply adding an L2 penalty on the unnormalised prediction scores (logits) of the network.

We evaluate the utility of spectral decoupling as an implicit bias mitigation method in the context of medical imaging. We use simulation experiments to show that spectral decoupling increases networks’ robustness to data distribution shifts and can be used to train generalisable networks on datasets with a strong superficial correlation. The findings are then evaluated by training prostate cancer and COVID-19 classifiers, where the networks trained with spectral decoupling achieve significantly higher performance on all evaluation datasets.

2 MATERIALS AND METHODS

2.1 Spectral decoupling

In spectral decoupling, the network is regularised by imposing an L2 penalty on the unnormalised outputs of the last layer of the network, or logits \hat{y} . This penalty avoids the conditions leading to gradient starvation in networks trained with cross-entropy loss. Two variants of the penalty are defined in the original paper [2].

$$\frac{\lambda}{2} \|\hat{y}\|_2^2 \quad (1)$$

$$\frac{\lambda}{2} \|\hat{y} - \gamma\|_2^2 \quad (2)$$

For Equation 1, there is a single tunable hyper-parameter λ . For Equation 2, hyper-parameters λ and γ are tuned separately for each class, a total of four hyper-parameters for the binary classification task in our study. For Equation 1, the tuned hyper-parameter is $\lambda = 0.01$ for the experiments in Sections 3.2 and 3.3. For Equation 2, the tuned hyper-parameters are $\lambda_{neg} = 0.0969$, $\gamma_{neg} = 1.83$, $\lambda_{pos} = 0.000698$ and $\gamma_{pos} = 2.61$ for the experiment in Section 3.1, and $\lambda_{neg} = 0.01$, $\gamma_{neg} = 0$, $\lambda_{pos} = 0.001$ and $\gamma_{pos} = 1$ for the experiment in Section 3.4.

2.2 Datasets

2.2.1 Prostate

A total of 30 prostate cancer patient cases are annotated for classification into cancerous and benign tissue. All patients have undergone a radical prostatectomy at the Helsinki University Hospital between 2014 and 2015. Each case contains 14 to 21 tissue section slides of the prostate. Tissue sections have a thickness of $4 \mu\text{m}$ and were stained with hematoxylin and eosin in a clinical-grade laboratory at the Helsinki University Hospital Diagnostic Center, Department of Pathology. Two different scanners are used to obtain images of the tissue section slides

at 20x magnification. Larger macro slides (whole-mount, 2x3 inch slides) are scanned with an Axio Scan Z.1 scanner (Zeiss, Oberkochen, Germany), and the normal size slides with a Panoramic Flash III 250 scanner (3DHitech, Budapest, Hungary). From the 30 patient cases, five are set aside for a test set and four are used as a validation set during training and hyper-parameter tuning. The test set is further divided based on the scanner used to obtain the images. Digital slide images are cut into tiles with 1024×1024 pixels and 20% overlap, resulting in 4.7 million tiles with 10% containing cancerous tissue.

For external validation, a freely available prostate cancer dataset is used, containing tissue section slides from patients who have undergone a radical prostatectomy at the Radboud University Medical Center between 2006 and 2011 [13, 14]. The dataset contains images with 2500×2500 pixels annotated by a uropathologist as either cancerous or benign. Images are scanned with a Panoramic Flash II 250 scanner (3DHitech, Budapest, Hungary) at 20x magnification but later reduced to 10x magnification. These images are cut into tiles with 512×512 pixels and 20% overlap, resulting in 5655 tiles with 45% containing cancerous tissue.

All slide images are cut and processed with HistoPrep [15] and albumentations [16] is used for image augmentations. A summary of the prostate datasets is presented in Table 1.

2.2.2 COVID-19

For COVID-19 detection, we use large open-access repositories of chest radiographs. COVIDx8 dataset is compiled from five different open-source repositories and contains radiographs from over 15,000 patient cases from at least 51 countries, with over 1500 COVID-19 positive patient cases [17, 18, 19, 20, 21]. BIMCV \pm dataset (iteration 2) contains 3033 positive and 2743 negative COVID-19 patient cases, and 9171 radiographs, after exclusions, collected from the multiple same medical centres during the same time period [22]. Only PA and upright AP radiographs [17] with windowing information were selected from the BIMCV \pm dataset. PadChest dataset contains over 67,000 COVID-19 negative patient cases, and 114,227 radiographs from a single medical centre in Valencia, Spain [23]. 19 corrupted images were excluded from the PadChest dataset.

COVIDx8 dataset is reserved as an external dataset, and two training datasets are compiled by using only the BIMCV \pm dataset and by adding the PadChest and BIMCV \pm datasets together. 5% of both training datasets are set aside for validation.

2.3 Simulation datasets

Two simulation experiments are used to more closely investigate the utility of spectral decoupling as an implicit bias mitigation method. For both experiments, the dataset from Helsinki University Hospital described in Section 2.2 is modified in specific ways.

2.3.1 Cutout dataset

A dominant feature present in a real-world dataset could be, for example, a biological marker, a certain cancer type or a scanner artefact. To represent these kinds of features, 16 cutouts of 8×8 pixels are added to the images (Figure 1).

Table 1: Prostate datasets

Centre	Scanner	Slides	Tiles	Training data
Helsinki University Hospital	Pannoramic Flash III 250	Normal	1.0 million	Section 3.1
	Axio Scan Z.1	Macro	3.7 million	Sections 3.2 & 3.3
Radboud University Medical Center	Pannoramic Flash II 250	Both	5655	–

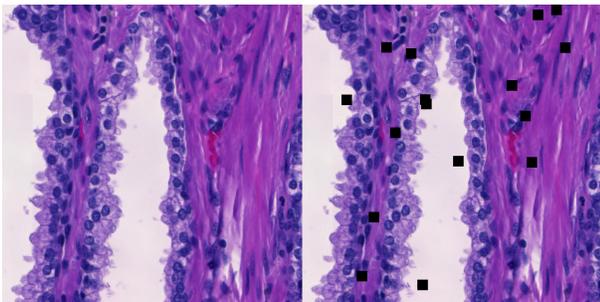


Figure 1: Left: Benign sample. Right: 16 cutouts of 8×8 pixels added to the benign sample.

For the experiment, 200,000 images are selected for the training set with equal amount of samples with cancerous and benign annotations. For the training set, cutouts are added to 25% and 2.5% of the benign and cancerous samples, respectively. This makes the presence of cutouts in the image spuriously correlated with a benign annotation. If the network overfits this correlation, cancerous samples with cutouts may be classified as benign. Thus for the test set, cutouts are added to all cancerous samples and none of the benign samples. For a control training set, cutouts are added to all images. Networks trained with this dataset provide a reference point of the performance with cutouts but without the spurious correlation.

2.3.2 Robustness dataset

Shifts from the training data distribution are common when evaluating the neural network with datasets from different medical centres. Small changes in the images due to differences in, for example, sample preparation or imaging equipment can cause shifts from the training data distribution. We assess the networks' robustness to these data distribution shifts, by applying transformations with increasing magnitudes to the images in the test set. Image sharpness was selected to represent one possible dataset shift for several reasons.

The `UniformAugment` augmentation strategy consists of applying random transformations with a uniformly sampled magnitude to the images before feeding them to the network [24]. Sharpening the image is included in the set of possible transformations [25], meaning that the network sees sharpened images during training. Thus, the data distribution shift caused by sharpening images is being explicitly mitigated, which should help the network to predict correct labels for evaluation images with higher sharpness. Blurring the image is not included in the set of possible transformations [25], meaning that the network will not see randomly blurred images during training. Thus, the data distribution shift caused by blurring the images will not be explicitly mitigated and the use of `UniformAugment` should not directly help the network with blurry evaluation images. By

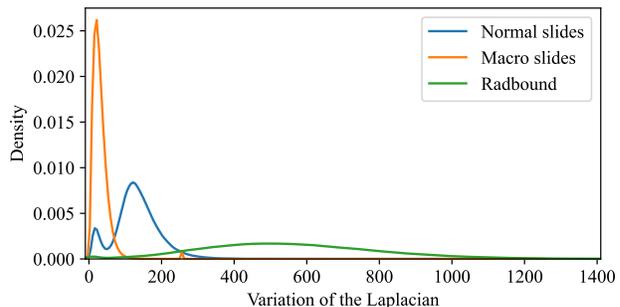


Figure 2: Kernel density estimation of the variance of the images after a Laplace transformation. A higher variance indicates a sharper image. The image is generated from the preprocessing metrics calculated by `HistoPrep` [15].

evaluating the network with increasingly sharpened or blurred images, it is possible to assess whether spectral decoupling can improve upon situations where the data distribution shift is, and is not explicitly addressed. Additionally, there are large differences in the sharpness values of real-world datasets from different medical centres and scanners (Figure 2).

Step-wise blurring is achieved by simple averaging with a $n \times n$ kernel, where $n \in \{2, \dots, 20\}$. Sharpened version of the image x_{sharp} is created by applying kernel

$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 9 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

to the original image x_{original} . Sharpness is then gradually increased by creating a new image x_{blend} with

$$x_{\text{blend}} = (1 - \alpha)x_{\text{original}} + \alpha x_{\text{sharp}},$$

where $\alpha \in \{0, 0.1, \dots, 1\}$ defines the amount of sharpness increase.

2.4 Training details

EfficientNet-b0 network [26], with dropout [27] and stochastic depth [28] of 20% and an input size of 224×224 , is used as a prostate cancer classifier for all experiments. For augmentation, the input images are randomly cropped and flipped, resized, and then transformed with `UniformAugment` [24], using a maximum of two transformations.

For COVID-19 detection, we replicate the training regimen from [6], where a DenseNet-121 network is pre-trained with the ImageNet dataset and then fine-tuned for 30 epochs as a binary COVID-19 classifier. All hyper-parameters, other than spectral decoupling, are set to values reported in the paper.

Table 2: Results of the simulation study with the cutout dataset on dominant features. The mean and standard deviation (SD) values are reported for each set of five trained networks.

Name	Accuracy (SD)	Recall (SD)
Weight decay	0.752 (0.019)	0.523 (0.039)
Spectral decoupling	0.837 (0.020)	0.715 (0.046)
Control + weight decay	0.875 (0.009)	0.832 (0.036)

EfficientNet-b0 networks are trained until convergence of the validation loss with either weight decay or spectral decoupling, keeping all other hyper-parameters equal. For spectral decoupling, Equation 2 is used for COVID-19 detection and the first simulation experiment on dominant features (Sections 3.4 and 3.1). Equation 1 is used for all other experiments (Sections 3.2 and 2.2). Bayesian optimisation is used to tune hyper-parameters for the experiment in Section 3.1. For all other experiments, a simple grid search over limited search spaces is used. Search spaces are defined as $S_1 = \{0.1, 0.01, \dots, 0.000001\}$ and $S_2 = \{-1, 0, 1, 2\}$, where $\lambda, \lambda_{pos}, \lambda_{neg} \in S_1$ and $\gamma_{pos}, \gamma_{neg} \in S_2$. Hyper-parameter tuning is done on the validation split for Equation 1, and on the test split for Equation 2.

Each experiment is repeated five times and the summary metrics for these runs are reported. All reported performance metrics are balanced between the classes when necessary and a cut-off value of 0.5 is used to obtain a binary label from the normalised predictions of the network. To compare paired receiver under the operating characteristic (ROC) curves, we use one-tailed DeLong’s test and report the Z-values and p-values [29]. Statistical significance is defined as a p-value less than 0.05.

3 EXPERIMENTS

In this section, the utility of using spectral decoupling as an implicit bias mitigation method is explored with both simulation and real-world experiments.

3.1 Dominant features

To assess the utility of spectral decoupling in situations where the training dataset contains a strong dominant feature, the cutout dataset defined in Section 2.3.1 is used. Five networks are trained with either spectral decoupling or weight decay on the training set. Additionally, five networks are trained on the control dataset with weight decay to provide a reference point of the performance under no spurious correlation caused by the dominant feature. The mean and standard deviation of the accuracy and recall metrics on the test data are reported in Table 2. Accuracy is defined as the fraction of all instances that were correctly identified, and recall as the fraction of positive instances that were correctly identified.

The use of spectral decoupling increases the accuracy by 8.5 percentage points over weight decay and almost reaches the performance of the network trained on the control dataset. The networks trained without spectral decoupling appear to make false predictions based on the dominant feature. As hyper-parameters were tuned on the test set, the results should be interpreted only

as a demonstration that spectral decoupling offers an important level of control over the features that are learned.

The simpler variant of spectral decoupling in Equation 1 did not increase the networks’ performance in any way, and only after tedious hyper-parameter tuning the variant in Equation 2 produced the reported results. Similar results were also reported with the real-world example in the original paper [2]. As hyper-parameter tuning can deter researchers from applying the method, we limit hyper-parameter tuning to a simple grid search over limited search spaces for all other experiments, as described in Section 2.4.

3.2 Robustness

To assess whether spectral decoupling increases neural networks’ robustness to data distribution shifts, five networks are trained with either spectral decoupling or weight decay and evaluated on the robustness dataset described in Section 2.3.2. Additionally, five networks are trained with weight decay but without `UniformAugment` to assess how much the augmentation strategy improves robustness. The robustness to data distribution shifts caused by sharpening and blurring the evaluation images are presented in Figure 3.

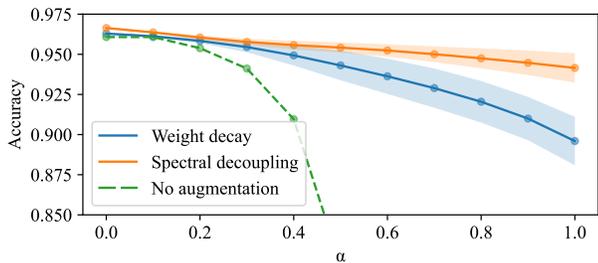
Performance of all networks trained with weight decay and without the augmentation strategy degrade to roughly 50% accuracy. Training the networks again with `UniformAugment` increases robustness to both data distribution shifts substantially. When the data distribution shift is included as a possible augmentation, the increase in accuracy is almost 40 percentage points with the most severe distribution shift (Figure 3a). When the data distribution shift is not included as a possible transformation, robustness is more similar with and without augmentation. This result demonstrates the importance of using augmentation as an explicit bias mitigation method.

Networks trained with spectral decoupling increase the robustness to data distribution shifts in both scenarios. Although the use of augmentation already increased the accuracy by almost 40 percentage points, the use of spectral decoupling is able to improve the accuracy by a further 4.6 percentage points with the most severe data distribution shift. The increase in accuracy is more pronounced with blurring, 12.4 percentage points with $n = 19$, which is not included as a possible transformation. These results show that spectral decoupling is able to complement and improve upon augmentation, as well as improve robustness to data distribution shifts that covered by augmentation.

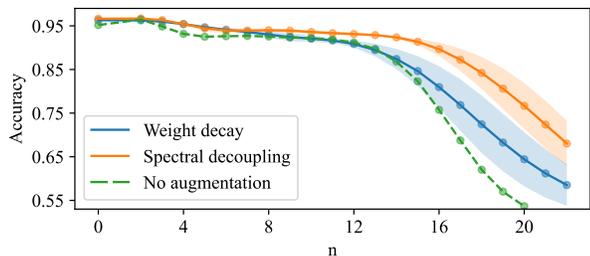
3.3 Prostate cancer detection

To assess whether the results of the simulation experiments translate into improvements in real-world datasets, we train networks with and without spectral decoupling to detect prostate cancer on haematoxylin and eosin stained whole slide images of the prostate. These networks are then evaluated on three different datasets described in Section 2.2.

The results are presented in Figure 4. Networks trained with spectral decoupling show higher performance on all evaluation datasets. The difference between weight decay and spectral decoupling gets more pronounced as we move further away from the training dataset distribution. Finally, there is a 10.6 percent-



(a) Sharpening.



(b) Blurring.

Figure 3: Robustness for data distribution shifts from the training data. The lines show the mean accuracy and the shaded regions represent one standard deviation around the mean.

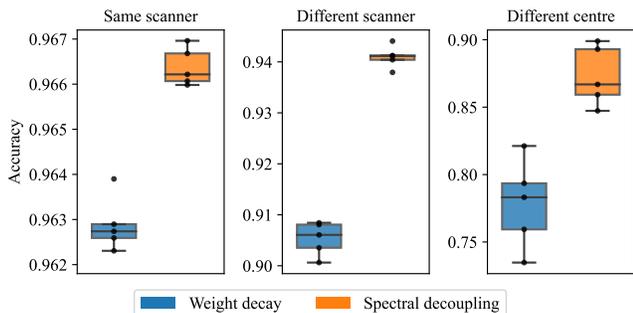


Figure 4: Neural network performance on evaluation datasets. Evaluation datasets move further from the training data distribution from left to right.

age point increase in accuracy over weight decay on the dataset from a different medical centre. The reported performances are not comparable between evaluation datasets, as each dataset has been annotated with a different strategy and thus contain different amounts of label noise.

3.4 COVID-19 detection

To assess whether spectral decoupling can help in real-world situations with strong dominant features and spurious correlations, we train 5 networks with and without spectral decoupling to detect COVID-19 positive patients in chest radiographs. Two different training datasets are used to train the networks and all networks are evaluated on the same external validation set, described in Section 2.2.2. We first train neural networks with the BIMCV± dataset, which represents an ideal situation where both the positive and negative samples originate from similar sources. Second, we train networks with the combined PadChest

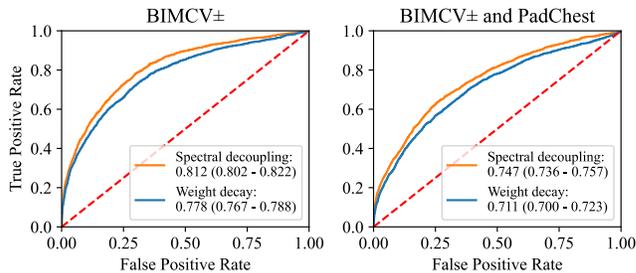


Figure 5: Receiver operating characteristic curves for COVID-19 detection. Inset values indicate the areas under these curves values and bootstrapped 95% confidence intervals.

and BIMCV± dataset. This dataset represents a situation where the network can easily achieve high performance by only learning to detect where a sample originates as most of the negative samples come from a single medical centre.

After training all networks, the predictions from each network are averaged to obtain ensemble predictions for both weight decay and spectral decoupling. ROC curves for ensemble predictions are presented in Figure 5, with bootstrapped ($n = 1000$) 95% confidence intervals (CI) for each area under the ROC curve (AUROC) value. Networks trained with spectral decoupling achieve significantly higher AUROC values for both BIMCV± (DeLong’s test: $Z = -15.914$, $-\log_{10}(p) = 56.597$) and the combined PadChest and BIMCV± (DeLong’s test: $Z = -13.553$, $-\log_{10}(p) = 41.418$) training datasets. On the BIMCV± dataset, weight decay and spectral decoupling achieve AUROCs of 0.812 (95% CI: 0.802 – 0.822) and 0.778 (95% CI: 0.767 – 0.788), respectively. With the combined PadChest and BIMCV± weight decay and spectral decoupling achieve AUROCs of 0.747 (95% CI: 0.736 – 0.757) and 0.711 (95% CI: 0.700 – 0.723), respectively.

When training networks with the combined PadChest and BIMCV± dataset, AUROC values of networks trained with either method decrease, although the number of training samples is increased over tenfold. The decrease in AUROC is similar for both weight decay and spectral decoupling, 0.065 and 0.067, respectively. This indicates that spectral decoupling is unable to mitigate bias in the combined dataset. As most of the negative samples originate from a single medical centre, shortcut learning seems to happen even though spectral decoupling encourages the network to learn more features. Detecting where a sample originates is especially easy with radiographs due to systematic differences between data repositories and medical centres, which could be exploited by a neural network [6]. Thus, the higher AUROC value of spectral decoupling is more likely due to increased robustness to data distribution shifts than avoidance of shortcut learning.

4 DISCUSSION

Generalisation performance is defined as the main challenge standing in the way of true clinical adoption of a neural network [7]. Van der Laak *et al.* [7] argue that there is a need for public datasets which are truly representative of clinical practice. Although this is indeed important, we argue that training datasets,

no matter how large, will never account for all possible variation caused by differences in imaging equipment, sample preparation and patient populations. Thus, it is crucial to use explicit and implicit bias mitigation methods to train neural networks which are robust to these variations.

Two explicit methods of bias mitigation have been proposed for medical imaging. Augmentation of the training samples is crucial as it substantially increases robustness for distribution shifts from the training data caused by differences in imaging equipment or sample preparation (Figure 3, [8]). Despite this, recent papers can be found that use no augmentation strategies ([30, 31]). Normalization of all images to a common standard would substantially reduce the distribution shifts [9, 10], but comes with a considerable computational cost. Both methods address important problems and should be complementary to any implicit methods of bias control.

Spectral decoupling is, to our knowledge, the first implicit bias mitigation method for addressing the generalisation issues in neural networks. The method is complementary to augmentation, increasing the robustness for distribution shifts already addressed with augmentation (Figure 3a). Above all, spectral decoupling significantly increases the robustness for distribution shifts not addressed by augmentation (Figure 3b).

The advantages of spectral decoupling can be clearly seen when the network is evaluated with out-of-distribution samples (Figures 3, 4 and 5). Neural networks trained with spectral decoupling retain their performance with differing samples, which is exactly what is required from neural network intended for clinical practice [7].

By encouraging the neural network to learn more features, spectral decoupling can also help in situations where the training dataset contains strong dominant features or spurious correlations (Table 2). This is crucial as the dominant features can also be inherent to the data, such as different cancer types. For example, with prostate cancer, different Gleason grades [32] are often unbalanced in the training set. Due to gradient starvation [11], the features of the underrepresented Gleason grades may not be learned by the neural network. Balancing the dataset, so that all Gleason grades are represented equally, is not easy or even desired as the grading is based on a continuous range of histological patterns.

In COVID-19 detection, the networks' performance decreased similarly for both weight decay and spectral decoupling (Figure 5), when training the networks on the combined BIMCV \pm and PadChest dataset. Radiographs contain systematic differences between data repositories and medical centres, such as laterality tokens and differences in the radiopacity of the image borders, which could arise from variations in patient position, radiographic projection or image processing [6]. These differences can be leveraged by neural networks to easily detect where a single radiograph originates. We speculate that spectral decoupling was unable to prevent shortcut-learning due to the ease of shortcut learning in the combined PadChest and BIMCV \pm dataset. Thus, removal of any obvious superficial correlations from the training dataset is crucial as there seems to be a limit of how much spectral decoupling can help with dominating features and spurious correlations.

5 CONCLUSIONS

Spectral decoupling is the first implicit bias mitigation method for training neural networks to be used across multiple medical centres. The method adds no computational costs, is easy-to-use, and complements and improves upon explicit bias mitigation methods. Our results recommend the use of spectral decoupling in all neural networks intended for clinical use.

ACKNOWLEDGMENTS

This work was supported by Cancer Foundation Finland [grant numbers 304667, 191118], Jane and Aatos Erkko Foundation [grant number 290520], Academy of Finland [grant number 322675] and Hospital District of Helsinki and Uusimaa [grant numbers TYH2018214, TYH2018222, TYH2019235, TYH2019249]. The authors also wish to acknowledge CSC – IT Center for Science, Finland, for generous computational resources and Wouter Bulten for directing us to the external prostate cancer dataset.

ETHICS STATEMENT

The images of the tissue slides are applied in this study based on national legislation and a research permission from the Helsinki University Hospital (§105).

DECLARATION OF COMPETING INTERESTS

The authors have no interests to declare.

CREDIT AUTHORSHIP CONTRIBUTION STATEMENT

Joona Pohjonen: Conceptualization; Data curation; Formal analysis; Investigation; Methodology; Project administration; Software; Validation; Visualization; Roles/Writing - original draft. **Carolin Stürenberg:** Data curation; Writing - review & editing. **Antti Rannikko:** Funding acquisition; Resources; Supervision; Writing - review & editing. **Tuomas Mirtti:** Data curation; Funding acquisition; Resources; Supervision; Writing - review & editing. **Esa Pitkänen:** Supervision; Writing - review & editing.

REFERENCES

- [1] Xiaoxuan Liu, Livia Faes, Aditya U Kale, Siegfried K Wagner, Dun Jack Fu, Alice Bruynseels, Thushika Mahendiran, Gabriella Moraes, Mohith Shamdas, Christoph Kern, et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *The lancet digital health*, 1(6):e271–e297, 2019.
- [2] Mohammad Pezeshki, Sékou-Oumar Kaba, Yoshua Bengio, Aaron Courville, Doina Precup, and Guillaume Lajoie. Gradient starvation: A learning proclivity in neural networks. *arXiv preprint arXiv:2011.09468*, 2020.

- [3] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.
- [4] Julia K Winkler, Christine Fink, Ferdinand Toberer, Alexander Enk, Teresa Deinlein, Rainer Hofmann-Wellenhof, Luc Thomas, Aimilios Lallas, Andreas Blum, Wilhelm Stolz, et al. Association between surgical skin markings in dermoscopic images and diagnostic performance of a deep learning convolutional neural network for melanoma recognition. *JAMA dermatology*, 155(10):1135–1141, 2019.
- [5] John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric Karl Oermann. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS medicine*, 15(11):e1002683, 2018.
- [6] Alex J DeGrave, Joseph D Janizek, and Su-In Lee. Ai for radiographic covid-19 detection selects shortcuts over signal. *Nature Machine Intelligence*, pages 1–10, 2021.
- [7] Jeroen van der Laak, Geert Litjens, and Francesco Ciompi. Deep learning in histopathology: the path to the clinic. *Nature medicine*, 27(5):775–784, 2021.
- [8] David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen van der Laak. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical image analysis*, 58:101544, 2019.
- [9] Thomas de Bel, Meyke Hermsen, Jesper Kers, Jeroen van der Laak, and Geert Litjens. Stain-transforming cycle-consistent generative adversarial networks for improved segmentation of renal histopathology. In M. Jorge Cardoso, Aasa Feragen, Ben Glocker, Ender Konukoglu, Ipek Oguz, Gozde Unal, and Tom Vercauteren, editors, *Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning*, volume 102 of *Proceedings of Machine Learning Research*, pages 151–163. PMLR, 08–10 Jul 2019.
- [10] Andrew Janowczyk, Ajay Basavanahally, and Anant Madabhushi. Stain normalization using sparse autoencoders (stanosa): application to digital pathology. *Computerized Medical Imaging and Graphics*, 57:50–61, 2017.
- [11] Remi Tachet des Combes, Mohammad Pezeshki, Samira Shabani, Aaron Courville, and Yoshua Bengio. On the learning dynamics of deep neural networks. *arXiv preprint arXiv:1809.06848*, 2018.
- [12] Augustin Cauchy. Méthode générale pour la résolution des systèmes d’équations simultanées. *Comp. Rend. Sci. Paris*, 25(1847):536–538, 1847.
- [13] Wouter Bulten, Péter Bándi, Jeffrey Hoven, Rob van de Loo, Johannes Lotz, Nick Weiss, Jeroen van der Laak, Bram van Ginneken, Christina Hulsbergen-van de Kaa, and Geert Litjens. Peso: Prostate epithelium segmentation on h&e-stained prostatectomy whole slide images. <https://zenodo.org/record/1485967#.YEn0viORpQK>, 2018. Accessed: 05.03.2021.
- [14] Wouter Bulten, Péter Bándi, Jeffrey Hoven, Rob van de Loo, Johannes Lotz, Nick Weiss, Jeroen van der Laak, Bram van Ginneken, Christina Hulsbergen-van de Kaa, and Geert Litjens. Epithelium segmentation using deep learning in h&e-stained prostate specimens with immunohistochemistry as reference standard. *Scientific reports*, 9(1):1–10, 2019.
- [15] Joon Pohjonen. Histoprep: Preprocessing large medical images for machine learning made easy! <https://github.com/jopo666/HistoPrep>, 2021. Accessed: 05.07.2021.
- [16] Alexander Buslaev, Vladimir I. Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A. Kalinin. Albumentations: Fast and flexible image augmentations. *Information*, 11(2), 2020.
- [17] Joseph Paul Cohen, Paul Morrison, Lan Dao, Karsten Roth, Tim Q Duong, and Marzyeh Ghassemi. Covid-19 image data collection: Prospective predictions are the future. *arXiv 2006.11988*, 2020.
- [18] Linda Wang, Zhong Qiu Lin, and Alexander Wong. Covidnet: a tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific Reports*, 10(1):19549, Nov 2020.
- [19] Emily B Tsai, Scott Simpson, Matthew Lungren, Michelle Hershman, Leonid Roshkovan, Errol Colak, Bradley J Erickson, George Shih, Anouk Stein, Jayshree Kalpathy-Cramer, et al. The rsna international covid-19 open annotated radiology database (ricord). *Radiology*, page 203957, 2021.
- [20] Tawsifur Rahman, Amith Khandakar, Yazan Qiblawey, Anas Tahir, Serkan Kiranyaz, Saad Bin Abul Kashem, Mohammad Tariqul Islam, Somaya Al Maadeed, Susu M Zughair, Muhammad Salman Khan, et al. Exploring the effect of image enhancement techniques on covid-19 detection using chest x-ray images. *Computers in biology and medicine*, 132:104319, 2021.
- [21] Muhammad E. H. Chowdhury, Tawsifur Rahman, Amith Khandakar, Rashid Mazhar, Muhammad Abdul Kadir, Zaid Bin Mahbub, Khandakar Reajul Islam, Muhammad Salman Khan, Atif Iqbal, Nasser Al Emadi, Mamun Bin Ibne Reaz, and Mohammad Tariqul Islam. Can ai help in screening viral and covid-19 pneumonia? *IEEE Access*, 8:132665–132676, 2020.
- [22] Maria De La Iglesia Vayá, Jose Manuel Saborit, Joaquim Angel Montell, Antonio Pertusa, Aurelia Bustos, Miguel Cazorla, Joaquin Galant, Xavier Barber, Domingo Orozco-Beltrán, Francisco García-García, et al. Bimcv covid-19+: a large annotated dataset of rx and ct images from covid-19 patients. *arXiv preprint arXiv:2006.01174*, 2020.
- [23] Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vayá. Padchest: A large chest x-ray image dataset with multi-label annotated reports. *Medical image analysis*, 66:101797, 2020.
- [24] Tom Ching LingChen, Ava Khonsari, Amirreza Lashkari, Mina Rafi Nazari, Jaspreet Singh Sambee, and Mario A

- Nascimento. Uniformaugument: A search-free probabilistic data augmentation approach. *arXiv preprint arXiv:2003.14348*, 2020.
- [25] Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation policies from data, 2019.
- [26] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019.
- [27] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [28] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer, 2016.
- [29] Elizabeth R DeLong, David M DeLong, and Daniel L Clarke-Pearson. Comparing the areas under two or more correlated receiver operating characteristic curves: a non-parametric approach. *Biometrics*, pages 837–845, 1988.
- [30] Yuri Tolkach, Tilmann Dohmgörger, Marieta Toma, and Glen Kristiansen. High-accuracy prostate cancer pathology using deep learning. *Nature Machine Intelligence*, 2(7):411–418, 2020.
- [31] Zabit Hameed, Sofia Zahia, Begonya Garcia-Zapirain, José Javier Aguirre, and Ana María Vanegas. Breast cancer histopathology image classification using an ensemble of deep learning models. *Sensors*, 20(16):4373, 2020.
- [32] Jonathan I Epstein, Lars Egevad, Mahul B Amin, Brett Delahunt, John R Srigley, and Peter A Humphrey. The 2014 international society of urological pathology (isup) consensus conference on gleason grading of prostatic carcinoma. *The American journal of surgical pathology*, 40(2):244–252, 2016.