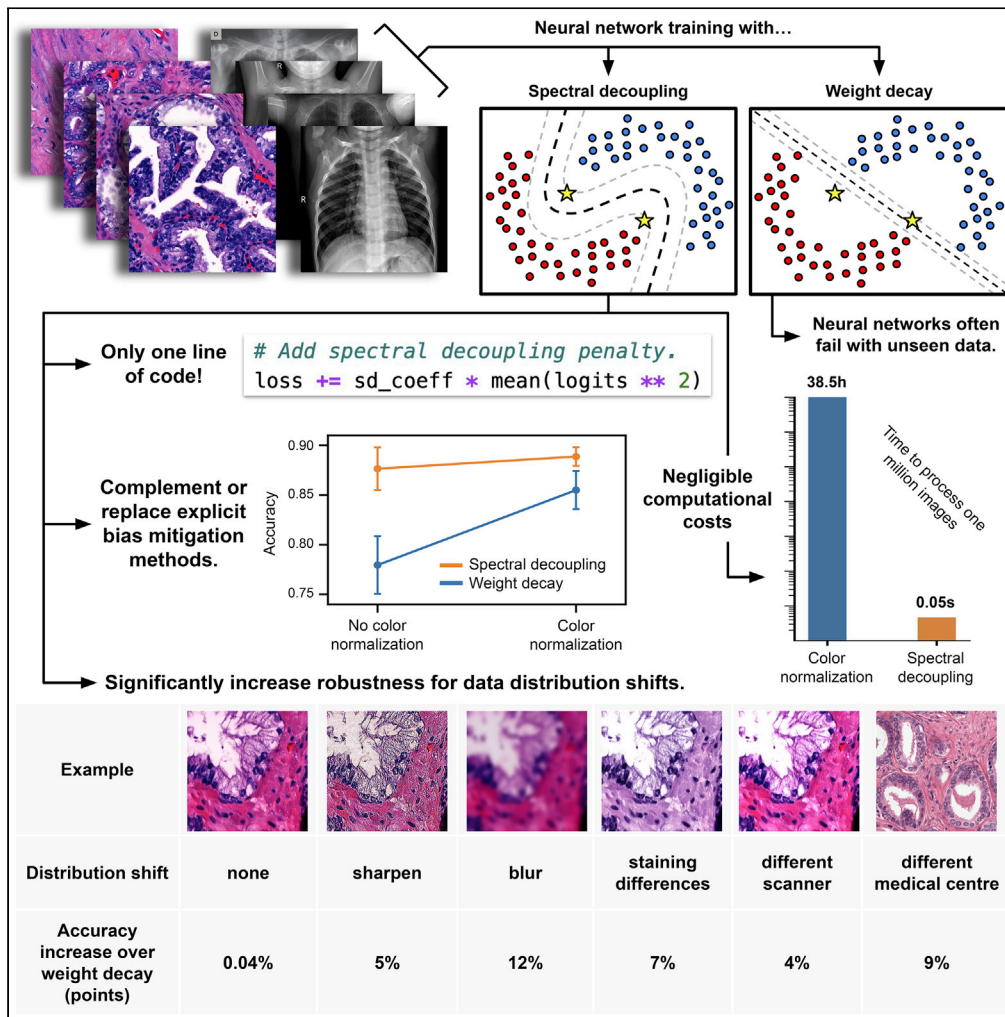


Article

Spectral decoupling for training transferable neural networks in medical imaging



Joona Pohjonen, Carolin Stürenberg, Antti Rannikko, Tuomas Mirtti, Esa Pitkänen

joona.pohjonen@helsinki.fi (J.P.)
esa.pitkanen@helsinki.fi (E.P.)

Highlights

We evaluate the first implicit bias mitigation method in medical imaging

Spectral decoupling increases robustness for distribution shifts and shortcut learning

Complement or replace explicit mitigation methods, such as color normalization

Up to 9.5% point higher accuracy on external datasets with one line of code



Article

Spectral decoupling for training transferable neural networks in medical imaging

Joonas Pohjonen,^{1,7,*} Carolin Stürenberg,¹ Antti Rannikko,^{1,2} Tuomas Mirtti,^{1,3,6} and Esa Pitkänen^{4,5,6,*}

SUMMARY

Many neural networks for medical imaging generalize poorly to data unseen during training. Such behavior can be caused by overfitting easy-to-learn features while disregarding other potentially informative features. A recent implicit bias mitigation technique called spectral decoupling provably encourages neural networks to learn more features by regularizing the networks' unnormalized prediction scores with an L2 penalty. We show that spectral decoupling increases the networks' robustness for data distribution shifts and prevents overfitting on easy-to-learn features in medical images. To validate our findings, we train networks with and without spectral decoupling to detect prostate cancer on tissue slides and COVID-19 in chest radiographs. Networks trained with spectral decoupling achieve up to 9.5 percent point higher performance on external datasets. Spectral decoupling alleviates generalization issues associated with neural networks and can be used to complement or replace computationally expensive explicit bias mitigation methods, such as stain normalization in histological images.

INTRODUCTION

Neural networks have been adapted to many medical imaging tasks with impressive results, often surpassing human counterparts in consistency, speed, and accuracy (Liu et al., 2019). However, these networks are prone to overfit easy-to-learn or statistically dominant features, while disregarding other potentially informative features. This leads to poor generalization to data generated by different medical centers, reliance on the dominant features, and lack of robustness (Geirhos et al., 2020; Pezeshki et al., 2020). 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 (Winkler et al., 2019). Owing 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 (Zech et al., 2018) and recent neural networks for COVID-19 detection rely on confounding factors rather than actual medical pathology (DeGrave et al., 2021). Even small differences in the sharpness of images from two different scanners can degrade the performance of neural networks significantly (see **Robustness** section).

Although generalization issues need to be solved before any neural networks can be applied in clinical practice, the phenomenon is still poorly understood (van der Laak et al., 2021). This may be because the detection of generalization issues is hard and often requires state-of-the-art methods of explainable AI (DeGrave et al., 2021). An external dataset is one of the only methods of testing generalization performance, although it will uncover generalization issues only when the neural network fails to generalize to the dataset. If a neural network achieves high overall accuracy on the external dataset, it may still always fail for some subset of samples. Any particular external dataset may also contain the same sources of bias as the training data.

Explicit methods have been proposed to address specific sources of bias, like using augmentation to address staining differences in tissue section slides (Tellez et al., 2019) or normalizing each image with a common standard (de Bel et al., 2019; Janowczyk et al., 2017). 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

¹Research Program in Systems Oncology, University of Helsinki, Helsinki, Finland

²Department of Urology, Helsinki University Hospital, Helsinki, Finland

³Department of Pathology, Helsinki University Hospital, Helsinki, Finland

⁴Institute for Molecular Medicine Finland (FIMM), HiLIFE, University of Helsinki, Helsinki, Finland

⁵Research Program in Applied Tumor Genomics, University of Helsinki, Helsinki, Finland

⁶Senior author

⁷Lead contact

*Correspondence: joona.pohjonen@helsinki.fi (J.P.), esa.pitkanen@helsinki.fi (E.P.)
<https://doi.org/10.1016/j.isci.2022.103767>



Table 1. Results of the simulation study with the cutout dataset on dominant features

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)

The mean and SD (SD) values are reported for each set of five trained networks.

generalization issues (Geirhos et al., 2020). Shortcut-learning occurs mainly because of gradient starvation, where gradient descent updates the parameters of a neural network in directions capturing only dominant features, thus starving the gradient from other features (des Combes et al., 2018). The gradient descent algorithm finds a local optimum by taking small steps toward the opposite sign of the derivative, the direction of the steepest descent (Cauchy, 1847). The recently proposed method of spectral decoupling (Pezeshki et al., 2020) 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 unnormalized 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 generalizable 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.

RESULTS

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

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 [Simulation datasets](#) is used. Five networks are trained with either spectral decoupling or weight decay on the training set. In addition, 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 SD of the accuracy and recall metrics on the test data are reported in [Table 1](#). 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, although the class activation maps (Chattopadhyay et al., 2018) of the trained neural networks, do not significantly differ between weight decay and spectral decoupling. As hyper-parameters were tuned on the test set, the results should be interpreted only as a demonstration that spectral decoupling can offer 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 extensive hyper-parameter tuning, [Equation 2](#) produced the reported results. The hyper-parameter tuning was sensitive to the selected parameters, and even small changes to the final values significantly reduced the accuracy of the neural network. Similar results were also reported with the real-world example in the original paper (Pezeshki et al., 2020). As extensive 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 [Spectral decoupling](#).

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 [Simulation datasets](#). In addition, five networks are trained with weight decay but

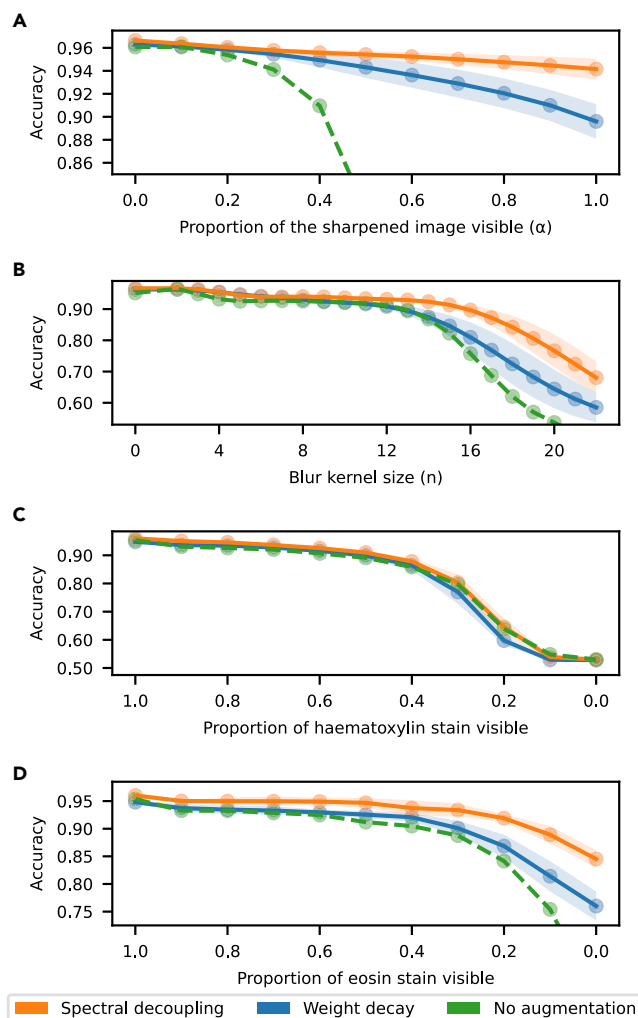


Figure 1. Robustness for data distribution shifts from the training data

The lines show the mean accuracy and the shaded regions represent one SD around the mean.

without UniformAugment to assess how much the augmentation strategy improves robustness. The robustness to data distribution shifts caused by sharpening, blurring, and reducing the intensity of either hematoxylin or eosin stain are presented in Figure 1.

Performance of all networks trained with weight decay and without the augmentation strategy degrades to roughly 50% accuracy. Training the networks again with UniformAugment significantly increases robustness to all data distribution shifts except with hematoxylin stain intensity reduction (Figure 1C). When the data distribution shift is included as a possible augmentation (Figure 1A), the increase in accuracy is almost 40 percentage points with the most severe distribution shift. When the data distribution shift is not included as a possible transformation (Figures 1B–D), robustness is more similar with and without augmentation. This result demonstrates the importance of using augmentation as an explicit bias mitigation method.

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 (Figure 1A). The increase in accuracy is more pronounced with blurring, 12.4 percentage points with $n=19$ (Figure 1B), and eosin stain intensity reduction, where networks trained with spectral decoupling achieve 1.2 to 8.5 percentage points higher accuracy with a 0.9 to 0.0 multiplier (Figure 1D). These data distribution shifts are not included as possible transformations in UniformAugment,

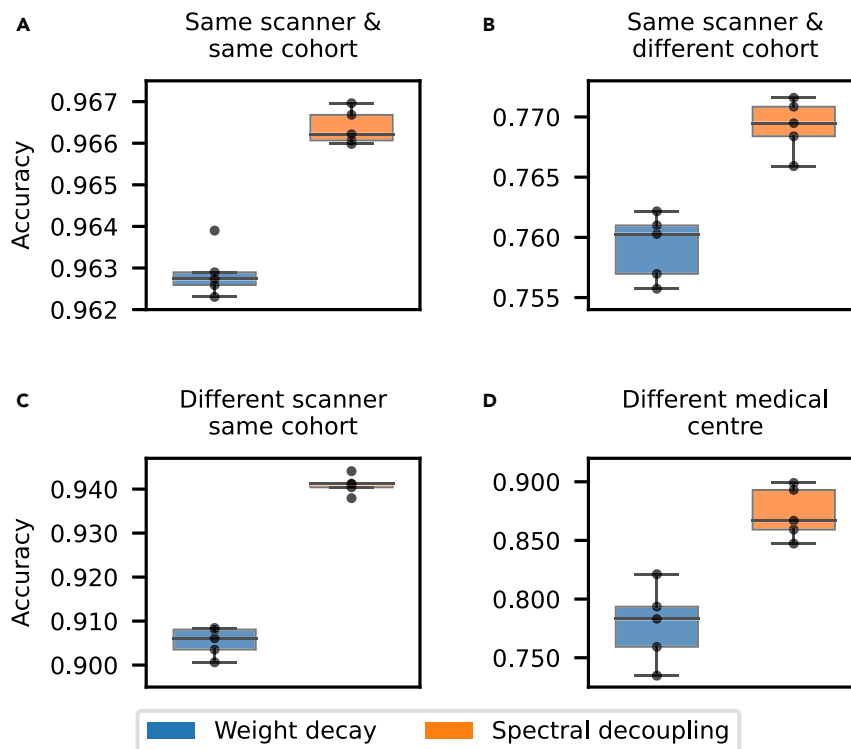


Figure 2. Neural network performance on evaluation datasets

(A–D) Each consecutive evaluation dataset moves further from the training data distribution. Networks trained with spectral decoupling improve accuracy by 0.35 (A), 1.0 (B), 3.6 (C) and 9.5 (D) percentage points over weight decay. All networks are trained with UniformAugment.

and thus not explicitly controlled. With hematoxylin stain intensity reduction, all networks degrade similarly in performance (Figure 1C). These results show that spectral decoupling is able to significantly complement and improve upon augmentation, as well as improve robustness to data distribution shifts that are not explicitly controlled by augmentation.

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 H&E stained whole slide images of the prostate. These networks are then evaluated on four different datasets described in Prostate dataset.

The results are presented in Figure 2. 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 9.5 percentage point increase in accuracy over weight decay on the dataset from a different medical center. 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.

To further explore why networks trained without spectral decoupling fail to generalize to the dataset from Radboud University Medical Center (Figure 2D), the robustness to H&E stain intensities are explored in Figures 3A and 3B. Spectral decoupling is less sensitive to both H&E stain intensity reduction and interestingly, networks trained with weight decay actually increase in accuracy when reducing the eosin stain intensity. This indicates that the difference between spectral decoupling and weight decay performance in Figure 2D, may be partly because of differences in the stain intensities between the two medical centers. To explore this possibility, the stain intensities of the external dataset are normalized with the Macenko method (Macenko et al., 2009) to match the training data stain intensities and the resulting performance increases are

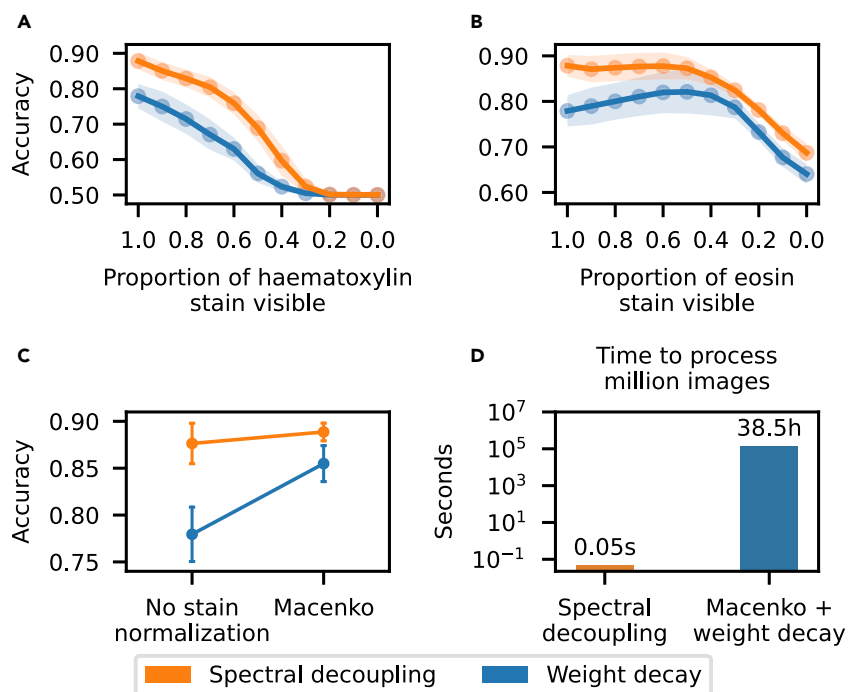


Figure 3. Spectral decoupling can complement or even replace computationally heavy stain normalization methods

(A and B) Robustness to data distribution shifts, on the external dataset, caused by heematoxylin (A) or eosin (B) stain intensity reduction.

(C) Network accuracy increases when normalizing H&E stain intensities with the Macenko method.

(D) Comparison of the computational requirements between spectral decoupling and the Macenko method. Images per seconds estimation for spectral decoupling is calculated with a Equation 1, where \hat{y} is a 512×1 matrix and Macenko stain normalization is performed on resized images of size 224×224 .

reported in Figure 3C. Both networks trained with either spectral decoupling or weight decay benefit from stain normalization. Stain normalization is especially beneficial for networks trained with weight decay, where the mean network accuracy is increased by 7.5 percentage points. Networks trained with spectral decoupling still perform better than networks trained with weight decay coupled with stain normalization. These results demonstrate that spectral decoupling can complement or even replace normalization methods, with negligible computational requirements (Figure 3D).

COVID-19 detection

To assess whether spectral decoupling can help in real-world situations with strong dominant features and spurious correlations, we train five 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 COVID-19 dataset. We first train neural networks with the BIMCV \pm 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 and BIMCV \pm 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 center.

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 4, with bootstrapped ($n = 1,000$) 95% CIs (CI) for each area under the ROC curve (AUROC) value. Networks trained with spectral decoupling achieve significantly higher AUROC values for both BIMCV \pm (De-Long's test: $Z = -15.914, p = 10^{-56}$) and the combined PadChest and BIMCV \pm (De-Long's test: $Z = -13.553, p = 10^{-41}$) training datasets. On the BIMCV \pm dataset, weight decay and spectral decoupling achieve AUROCs

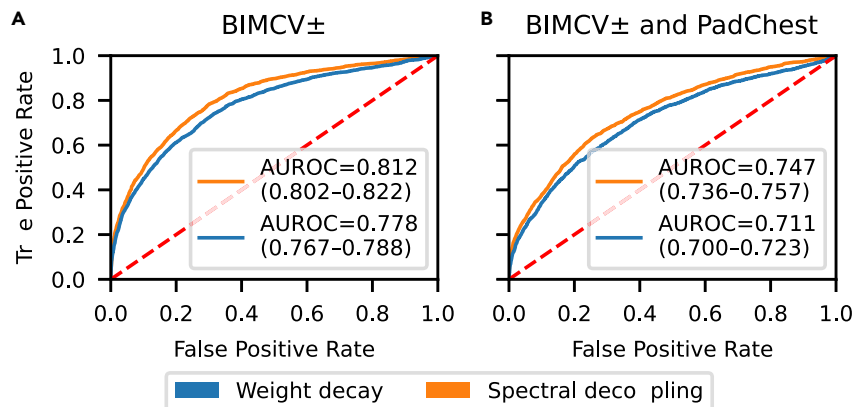


Figure 4. Receiver operating characteristic (ROC) curves for COVID-19 detection

Inset values indicate the areas under the ROC (AUROC) values and bootstrapped 95% CIs. Networks trained with spectral decoupling achieve significantly higher AUROC values compared to networks trained with weight decay.

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 10-fold. The decrease in AUROC is similar for 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 center, 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 because of systematic differences between data repositories and medical centers, which could be exploited by a neural network (DeGrave et al., 2021). Thus, the higher AUROC value of spectral decoupling is more likely because of increased robustness to data distribution shifts than avoidance of shortcut learning.

DISCUSSION

Generalization performance is defined as the main challenge standing in the way of true clinical adoption of a neural network (van der Laak et al., 2021). Van Der Laak et al. (2021) 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 variations caused by differences in imaging equipment, sample preparation, and patient populations. Thus, it is crucial to couple extensive multisource datasets with explicit and implicit bias mitigation methods to train neural networks which are robust to unseen 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 1, Tellez et al., 2019). Thus, it is strongly recommended to use extensive augmentation strategies for training neural networks intended for clinical practice. Normalization of all images to a common standard would substantially reduce the distribution shifts (de Bel et al., 2019; Janowczyk et al., 2017; Swiderska-Chadaj et al., 2020), but comes with a considerable computational cost (Figure 3D). 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 generalization issues in neural networks. The method is complementary to augmentation, increasing the robustness for distribution shifts already addressed with augmentation (Figure 1A). Above all, spectral decoupling significantly increases the robustness for distribution shifts not addressed by augmentation (Figure 1B) and could be used to replace computationally expensive stain normalization methods (Figure 3C).

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 1). 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 (Epstein et al., 2016) are often unbalanced in the training set. Owing to gradient starvation (des Combes et al., 2018), 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 4), when training the networks on the combined BIMCV \pm and the PadChest dataset. Radiographs contain systematic differences between data repositories and medical centers, 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 (DeGrave et al., 2021). These differences can be easily leveraged by neural networks to detect where a single radiograph originates. We speculate that spectral decoupling was unable to prevent shortcut-learning because of the ease of shortcut learning in the combined PadChest and BIMCV \pm dataset. In addition, our results showing the ability to prevent shortcut learning (Table 1) were obtained after considerable hyper-parameter optimization and no significant differences could be seen in the class activation maps between networks trained with either weight decay or spectral decoupling. 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.

The advantages of spectral decoupling can be clearly seen when the network is evaluated with out-of-distribution samples (Figures 1, 2, and 4). Neural networks trained with spectral decoupling retain their performance with samples further from the training data distribution, which is exactly what is required from neural networks intended for clinical practice (van der Laak et al., 2021). Although using an external dataset may not reveal all generalization problems, it is clear that without spectral decoupling the neural networks fail to generalize to this particular external dataset from Radboud University Medical Center (Figures 2D and 3). Even in COVID-19 detection, where spectral decoupling seems to fail in preventing shortcut learning, the performance of the network is significantly increased over the state-of-the-art.

Conclusions

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

Limitations of the study

Spectral decoupling is shown, by a simulation experiment, to offer an important level of control over the features that are learned in the 'dominant features' section. Despite this, spectral decoupling is unable to prevent shortcut learning as described in the COVID-19 detection section. We speculate this was because of the ease of shortcut learning in the training dataset, as mentioned in the discussion section. It is also possible spectral decoupling achieves significantly higher performance solely because of increased robustness to data distribution shifts and not also through the prevention of shortcut learning.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- METHOD DETAILS
 - Spectral decoupling
 - Prostate dataset

- COVID-19 dataset
- Simulation datasets
- Cutout simulation dataset
- Robustness simulation dataset
- **QUANTIFICATION AND STATISTICAL ANALYSIS**

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2022.103767>.

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. 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).

AUTHOR CONTRIBUTIONS

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

DECLARATION OF INTERESTS

The authors have no interests to declare.

Received: August 9, 2021

Revised: December 14, 2021

Accepted: January 10, 2022

Published: February 18, 2022

REFERENCES

- Bulten, W., Bándi, P., Hoven, J., van de Laak, J., van Ginneken, B., Hulsbergan van de Kaa, C., and Litjens, G. (2018). Peso: prostate epithelium segmentation on H&E-stained prostatectomy whole slide images. <https://zenodo.org/record/1485967#.YEn0vi0RpQK>.
- Bulten, W., Bándi, P., Hoven, J., van de Laak, J., van Ginneken, B., Hulsbergan van de Kaa, C., and Litjens, G. (2019). Epithelium segmentation using deep learning in H&E-stained prostate specimens with immunohistochemistry as reference standard. *Sci. Rep.* 9, 1–10.
- Buslaev, A., Iglovikov, V.I., Khvedchenya, E., Parinov, A., Druzhinin, M., and Kalinin, A.A. (2020). Albumentations: fast and flexible image augmentations. *Information* 11, 125. <https://doi.org/10.3390/info11020125>.
- Bustos, A., Pertusa, A., Salinas, J.-M., and de la Iglesia-Vayá, M. (2020). Padchest: a large chest x-ray image dataset with multi-label annotated reports. *Med. Image Anal.* 66, 101797.
- Cauchy, A. (1847). Méthode générale pour la résolution des systèmes d'équations simultanées. *Comp. Rend. Sci. Paris* 25, 536–538.
- Chattopadhyay, A., Sarkar, A., Howlader, P., and Balasubramanian, V.N. (2018). Grad-cam++: generalized gradient-based visual explanations for deep convolutional networks, pp. 839–847. <https://doi.org/10.1109/WACV.2018.00097>.
- Chowdhury, M.E.H., Rahman, T., Khandakar, A., Mazhar, R., Kadir, M.A., Mahbub, Z.B., Islam, K.R., Khan, M.S., Iqbal, A., Emadi, N.A., et al. (2020). Can AI help in screening viral and COVID-19 pneumonia? *IEEE Access* 8, 132665–132676. <https://doi.org/10.1109/ACCESS.2020.3010287>.
- Cohen, J.P., Morrison, P., Dao, L., Roth, K., Duong, T.Q., and Ghassemi, M. (2020). COVID-19 image data collection: prospective predictions are the future. *arXiv*, 2006.11988. <https://github.com/ieee8023/covid-chestxray-dataset>.
- Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q.V. (2019). Autoaugment: Learning Augmentation Policies from Data (*arXiv*), 1805.09501.
- de Bel, T., Hermsen, M., Kers, J., van der Laak, J., and Litjens, G. (2019). Stain-transforming cycle-consistent generative adversarial networks for improved segmentation of renal histopathology. In *Proceedings of the 2nd International Conference on Medical Imaging With Deep Learning*, M.J. Cardoso, A. Feragen, B. Glocker, E. Konukoglu, I. Ogu, G. Unal, and T. Vercauteren, eds. (PMLR), pp. 151–163. <http://proceedings.mlr.press/v102/de-bel19a.html>.
- De La Iglesia Vayá, M., Saborit, J.M., Montell, J.A., Pertusa, A., Bustos, A., Cazorla, M., Galant, J., Barber, X., Orozco-Beltrán, D., Garcia-Garcia, F., et al. (2020). BIMCV COVID-19+: a large annotated dataset of RX and CT images from COVID-19 patients. *arXiv*, preprint [arXiv:2006.01174](https://arxiv.org/abs/2006.01174).
- DeGrave, A.J., Janizek, J.D., and Lee, S.-I. (2021). AI for radiographic COVID-19 detection selects shortcuts over signal. *Nat. Machine Intelligence* 3, 1–10.

- DeLong, E.R., DeLong, D.M., and Clarke-Pearson, D.L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics* 44, 837–845.
- des Combes, R.T., Pezeshki, M., Shabanian, S., Courville, A., and Bengio, Y. (2018). On the learning dynamics of deep neural networks. *arXiv, preprint arXiv:1809.06848*.
- Epstein, J.I., Egevad, L., Amin, M.B., Delahunt, B., Strigley, J.R., and Humphrey, P.A. (2016). The 2014 international society of urological pathology (ISUP) consensus conference on gleason grading of prostatic carcinoma. *Am. J. Surg. Pathol.* 40, 244–252.
- Geirhos, R., Jacobsen, J., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., and Wichmann, F.A. (2020). Shortcut learning in deep neural networks. *Nat. Machine Intelligence* 2, 665–673.
- Huang, G., Liu, Z., van der Maaten, L., and Weinberger, K.Q. (2018). Densely Connected Convolutional Networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (IEEE)*.
- Huang, G., Sun, Y., Liu, Z., Sedra, D., and Weinberger, K.Q. (2016). Deep networks with stochastic depth. In *European Conference on Computer Vision (Springer International Publishing)*, pp. 646–661.
- Janowczyk, A., Basavanthally, A., and Madabhushi, A. (2017). Stain normalization using sparse autoencoders (stanosa): application to digital pathology. *Comput. Med. Imaging Graph.* 57, 50–61.
- LingChen, T.C., Khonsari, A., Lashkari, A., Nazari, M.R., Sambee, J.S., and Nascimento, M.A. (2020). Uniformaugmt: a search-free probabilistic data augmentation approach. *arXiv, preprint arXiv:2003.14348*.
- Liu, X., Faes, L., Kale, A.U., Wagner, S.K., Fu, D.J., Bruynseels, A., Mahendiran, T., Moraes, G., Shamdas, M., Kern, C., et al. (2019). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *Lancet Digital Health* 1, e271–e297.
- Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Guan, X., Schmitt, C., and Thomas, N.E. (2009). A method for normalizing histology slides for quantitative analysis, pp. 1107–1110. <https://doi.org/10.1109/ISBI.2009.5193250>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: an imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, eds. (Curran Associates, Inc.), pp. 8024–8035. <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Pezeshki, M., Kaba, S.-O., Bengio, Y., Courville, A., Precup, D., and Lajoie, G. (2020). Gradient starvation: a learning proclivity in neural networks. *arXiv, preprint arXiv:2011.09468*.
- Pohjonen, J. (2021). Histoprep: preprocessing large medical images for machine learning made easy!. <https://github.com/jop0666/HistoPrep>.
- Rahman, T., Khandakar, A., Qiblawey, Y., Tahir, A., Kiranyaz, S., Kashem, S.B.A., Islam, M.T., Al Maadeed, S., Zughair, S.M., Khan, M.S., et al. (2021). Exploring the effect of image enhancement techniques on COVID-19 detection using chest X-ray images. *Comput. Biol. Med.* 132, 104319.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Machine Learn. Res.* 15, 1929–1958.
- Swiderska-Chadaj, Z., de Bel, T., Blanchet, L., Baidoshvili, A., Vossen, D., van der Laak, J., and Litjens, G. (2020). Impact of rescanning and normalization on convolutional neural network performance in multi-center, wholeslide classification of prostate cancer. *Sci. Rep.* 10, 1–14.
- Tan, M., and Le, Q. (2019). Efficientnet: rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning (PMLR)*, pp. 6105–6114.
- Tellez, D., Litjens, G., Bándi, P., Bulten, W., Bokhorst, J.-M., Ciompi, F., and van der Laak, J. (2019). Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Med. Image Anal.* 58, 101544.
- Tsai, E.B., Simpson, S., Lungren, M., Hershman, M., Roshkovan, L., Colak, E., Erickson, B.J., Shih, G., Stein, A., Kalpathy-Cramer, J., et al. (2021). The RSNA International COVID-19 open annotated radiology database (RICORD). *Radiology* 299, 203957.
- van der Laak, J., Litjens, G., and Ciompi, F. (2021). Deep learning in histopathology: the path to the clinic. *Nat. Med.* 27, 775–784.
- Wang, L., Lin, Z.Q., and Wong, A. (2020). COVID-net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest x-ray images. *Sci. Rep.* 10, 19549. <https://doi.org/10.1038/s41598-020-76550-z>.
- Wightman, R. (2019). Pytorch Image Models. <https://doi.org/10.5281/zenodo.4414861>.
- Winkler, J.K., Fink, C., Toberer, F., Enk, A., Deinlein, T., Hofmann-Wellenhof, R., Thomas, L., Lallas, A., Blum, A., Stolz, W., et al. (2019). Association between surgical skin markings in dermoscopic images and diagnostic performance of a deep learning convolutional neural network for melanoma recognition. *JAMA Dermatol.* 155, 1135–1141.
- Zech, J.R., Badgeley, M.A., Liu, M., Costa, A.B., Titano, J.J., and Oermann, E.K. (2018). Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS Med.* 15, e1002683.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
PESO	Bulten et al. (2018, 2019)	N/A
COVIDx8	Cohen et al. (2020) , Wang et al. (2020) , Tsai et al. (2021) , Rahman et al. (2021) Chowdhury et al. (2020)	N/A
PadChest	Bustos et al. (2020)	N/A
BIMCV+/-	De La Iglesia Vayá et al. (2020)	N/A
Helsinki University Hospital (2014–2015)	This paper (not shared)	N/A
Helsinki University Hospital (2019–2020)	This paper (not shared)	N/A
Software and Algorithms		
PyTorch	https://pytorch.org	1.8
Albumentations	https://github.com/albumentations-team/albumentations	0.5.1
PyTorch image models	https://github.com/rwightman/pytorch-image-models	0.1.8
Python	https://www.python.org	3.8

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Joonas Pohjonen (joona.pohjonen@helsinki.fi).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- The patient data originating from the Helsinki University Hospital data reported in this study cannot be deposited in a public repository. This paper also used existing, publicly available data. These accession numbers for the datasets are listed in the [key resources table](#).
- This paper does not report original code. An interactive example demonstrating spectral decoupling with mock data has been published before ([Pezeshki et al., 2020](#)).
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

METHOD DETAILS

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 $\hat{\gamma}$, which is then added to cross-entropy loss, \mathcal{L}_{CE} . This penalty probably ([Pezeshki et al., 2020](#)) avoids the conditions leading to gradient starvation in networks trained with cross-entropy loss. Two variants of the penalty are defined as

$$\mathcal{L}_{CE} + \frac{\lambda}{2} \|\hat{\gamma}\|_2^2, \quad (\text{Equation 1})$$

$$\mathcal{L}_{CE} + \frac{\lambda}{2} \|\hat{\gamma} - \gamma\|_2^2. \quad (\text{Equation 2})$$

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. Pseudo-code for implementing Equation 1 is presented in Figure.

```
# net:      Neural network
# lambda_sd: Hyper-parameter for Equation 1.

for (images, targets) in loader:
    # Pass images through the network.
    logits = net(images)
    # Compute cross-entropy loss.
    loss = cross_entropy(logits, targets)
    # Add spectral decoupling penalty.
    loss += lambda_sd/2 * (logits**2).mean()
    # Optimization step.
    loss.backward()
    optimizer.step()
```

PyTorch style pseudocode for Equation 1.

A simple grid search is used to optimize the hyper-parameters in Sections 2.2, 2.3, and 2.4. Bayesian optimisation is used in Section 2.1. Search spaces for the grid search are defined as $S_1 = \{0.1, 0.01, \dots, 0.000001\}$, $S_2 = \{-1, 0, 1, 2\}$, where $\lambda, \lambda_{pos}, \lambda_{neg} \in S_1$ and $\gamma_{pos}, \gamma_{neg} \in S_2$. Hyper-parameter optimization is done on the validation split, except for Equation 2 in Section 2.1, where we perform optimization straight on the test split. For Equation 1, the tuned hyper-parameter is $\lambda = 0.01$. 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 2.1, and $\lambda_{neg} = 0.01$, $\gamma_{neg} = 0$, $\lambda_{pos} = 0.001$ and $\gamma_{pos} = 1$ for the experiment in Section 2.4.

Prostate dataset

A total of 30 prostate cancer patient cases are annotated for classification into cancerous and benign tissue, where the cancerous areas were annotated in consensus by two observers (C.S., T.M.). All patients have undergone 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 μm 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, 2 \times 3 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 \times 1024 pixels and 20% overlap, resulting in 4.7 million tiles with 10% containing cancerous tissue.

To test the differences between cohorts from the same medical centre, another set of 60 prostate cancer patient cases are annotated into cancerous and benign tissue by one of six experienced pathologists. All patients have undergone radical prostatectomy at the Helsinki University Hospital between 2019 and 2020. Each case contains 10 to 21 normal and macro tissue section slides of the prostate. Tissue sections have a thickness of 4 μm and are also stained with hematoxylin and eosin in a clinical-grade laboratory at the Helsinki University Hospital Diagnostic Center, Department of Pathology. All slides are scanned with an Axio Scan Z.1 scanner (Zeiss, Oberkochen, Germany). Digital slide images are cut into tiles with 1024 \times 1024 pixels and 20% overlap, resulting in 13.1 million tiles with 16% 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 (Bulten et al., 2018, 2019). The dataset contains images with 2500 \times 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 \times 512 pixels and 20% overlap, resulting in 5,655 tiles with 45% containing cancerous tissue.

All digital slide images are cut and processed with HistoPrep (Pohjonen, 2021). A summary of the prostate datasets is presented in Table.

Prostate datasets					
Center	Scanner	Slides	Tiles	Train data	Test data
Helsinki University Hospital	Pannoramic Flash III 250	Normal	1.0 million	Dominant features	Prostate cancer detection
	Axio Scan Z.1	Macro	3.7 million	Robustness and Prostate cancer detection	Robustness and Prostate cancer detection
	Axio Scan Z.1	Macro	13.1 million	–	Section Prostate cancer detection
Radboud University Medical Center	Pannoramic Flash II 250	Both	5,655	–	Section Prostate cancer detection

COVID-19 dataset

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 (Chowdhury et al., 2020; Cohen et al., 2020; Rahman et al., 2021; Tsai et al., 2021; Wang et al., 2020). BIMCV± dataset (iteration 2) contains 3,033 positive and 2,743 negative COVID-19 patient cases, and 9,171 radiographs, after exclusions, collected from the multiple same medical centres during the same time period (De La Iglesia Vayá et al., 2020). Only PA and upright AP radiographs (Cohen et al., 2020) with windowing information were selected from the BIMCV± dataset. PadChest dataset contains over 67,000 COVID-19 negative patient cases, and 114,227 radiographs from a single medical centre in Valencia, Spain (Bustos et al., 2020). 19 corrupted images were excluded from the PadChest dataset.

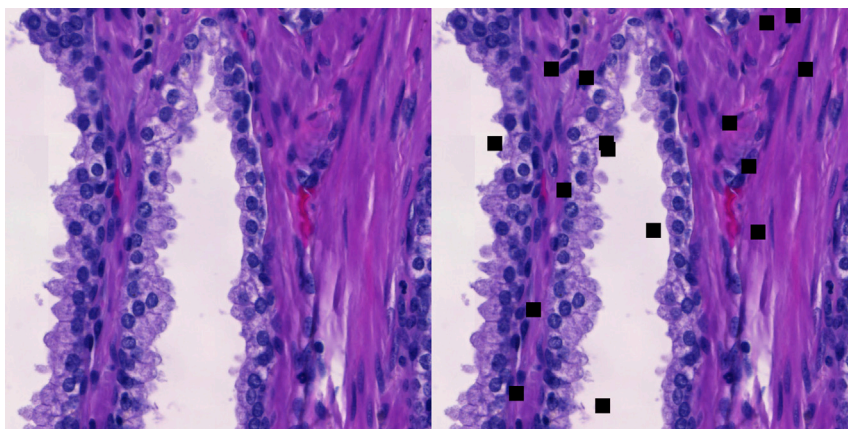
COVIDx8 dataset is reserved as an external dataset, and two training datasets are compiled by using only the BIMCV± dataset and by adding the PadChest and BIMCV± datasets together. 5% of both training datasets are set aside for validation.

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 9.2 is modified in specific ways.

Cutout simulation 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).



Example of the cutout operation

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

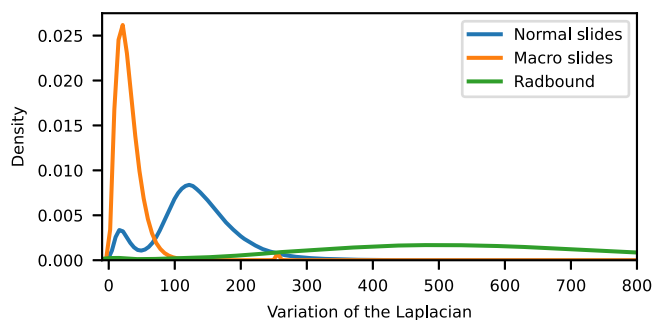
For the experiment, 200,000 images are selected for the training set with an equal amount of samples with cancerous and benign annotations. For the training set, cutouts are added to 25 and 2.5% of the benign and cancerous samples, respectively. This makes the presence of cutouts in the image spuriously 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.

Robustness simulation 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 and stain intensity were selected to represent possible dataset shifts caused by differences in the used scanner and sample preparation, respectively.

The UniformAugment augmentation strategy consists of applying random transformations with a uniformly sampled magnitude to the images before feeding them to the network (LingChen et al., 2020). Sharpening the image is included in the set of possible transformations (Cubuk et al., 2019), 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 (Cubuk et al., 2019), 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 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).



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 (Pohjonen, 2021).

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

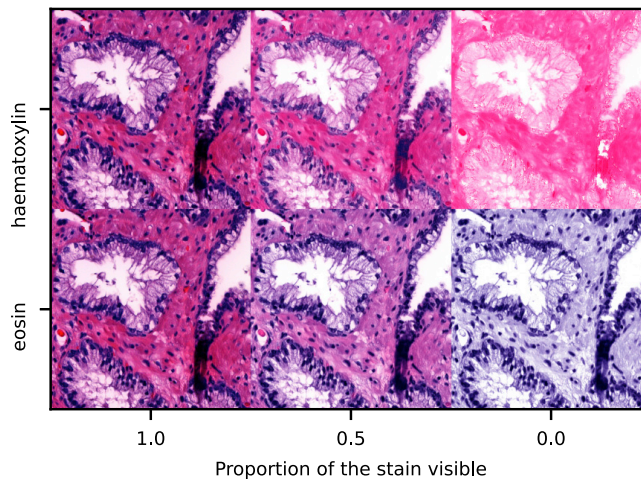
$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 9 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

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.

To assess the data distribution shifts caused by differences in sample preparation, the intensity of haematoxylin and eosin stains are computationally modified. Haematoxylin highlights cell nuclei, and eosin cytoplasm, connective tissue and muscle. The stain intensities depend on multiple steps in the staining process, and thus the final colour distribution of the slide images varies a lot (Tellez et al., 2019). The stain intensity modification is achieved by first separating the haematoxylin and eosin stains with the Macenko method (Macenko et al., 2009). The concentrations of each stain can then be reduced by multiplication with a value between 0 and 1 before the stains are combined back together. An example of the method is shown in Figure.



Separation of the hematoxylin and eosin stains with the Macenko method (Macenko et al., 2009).

QUANTIFICATION AND STATISTICAL ANALYSIS

EfficientNet-b0 network (Tan and Le, 2019), with dropout (Srivastava et al., 2014) and stochastic depth (Huang et al., 2016) 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 (LingChen et al., 2020), using a maximum of two transformations. Each network is trained for 90 epochs, with a learning rate of $0.005 \frac{\text{batch size}}{512}$ and cosine scheduling. Weight decay of 0.0001 is used for networks trained without spectral decoupling. When training neural networks with spectral decoupling, weight decay is disabled.

For COVID-19 detection, we replicate the training regimen from (DeGrave et al., 2021), where a DenseNet-121 network (Huang et al., 2018) 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. Training and validation curves for the trained networks are shown in Figure S1.

For spectral decoupling, Equation 2 is used for the first simulation experiment on dominant features (Section 2.1) and COVID-19 detection (Section 2.4). Equation 1 is used for all other experiments (Sections 2.2 and 2.3).

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 (DeLong et al., 1988).

PyTorch (version 1.8) (Paszke et al., 2019) is used for training the neural networks, timm (version 0.1.8) (Wightman, 2019) for building the neural networks and albumentations (version 0.5.1) (Buslaev et al., 2020) for image augmentations.