Synthetic Privileged Information Enhances Medical Image Representation Learning

Lucas Farndale[1,2] ✉ ⓘ, Chris Walsh[1,2] ⓘ, Robert Insall[1,2,3] ⓘ, and Ke Yuan[1,2] ⓘ

[1] University of Glasgow, Glasgow, Scotland, UK
[2] Cancer Research UK Scotland Institute, Glasgow, Scotland, UK
[3] University College London, London, England, UK
`lucas.farndale@glasgow.ac.uk`

Abstract. Multimodal self-supervised representation learning has consistently proven to be a highly effective method in medical image analysis, offering strong task performance and producing biologically informed insights. However, these methods heavily rely on large, paired datasets, which is prohibitive for their use in scenarios where paired data does not exist, or there is only a small amount available. In contrast, image generation methods can work well on very small datasets, and can find mappings between unpaired datasets, meaning an effectively unlimited amount of paired synthetic data can be generated. In this work, we demonstrate that representation learning can be significantly improved by synthetically generating paired information, both compared to training on either single-modality (up to $4.4\times$ error reduction) or authentic multi-modal paired datasets (up to $5.6\times$ error reduction).

Keywords: Image Generation · Pathology · Knowledge Distillation

1 Introduction

Recent work [1,2,3] has established the efficacy of integrating multiple sources of data to improve representation learning for histopathology. It has been demonstrated that integrating additional information, either advanced, high content information or something more routine, can greatly enhance the features that are learned by models, as it can act as an additional source of implicit supervision. This is particularly notable for privileged information – information that is available during training but not inference – which can be distilled into models to greatly improve single-modality downstream task performance where the privileged information is not available. Paired datasets can be used to distil knowledge from the privileged data to the primary data, producing much more biologically informative representations and improving performance on a variety of a priori unseen tasks during evaluation by only training a small classifier head [1].

A major drawback of these approaches, however, is the reliance on large datasets where data are paired one-to-one during training. In many cases, such as histology, the privileged data of interest only exist in small datasets or in distinct,
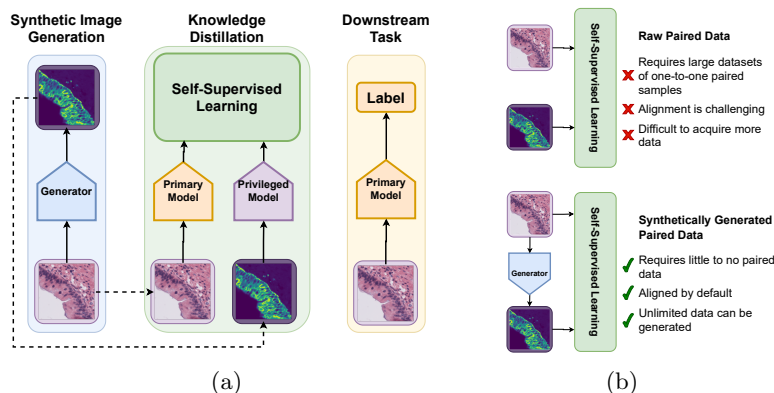
Fig. 1: (a) Schematic of synthetic paired data being generated and passed to the self-supervised model (b) Comparison of standard multimodal self-supervised approaches against our proposed method using synthetically generated data. Images from [4].

unpaired datasets which have not been previously amenable to these approaches. Image generation methods typically require orders of magnitude less data than self-supervised methods, with high quality image generation being possible with hundreds of examples (e.g. [4]), compared to hundreds of thousands for self-supervised learning [5]. Furthermore, cycleGAN [6] does not require a one-to-one pairing of images and can achieve excellent results for image generation. We may therefore construct arbitrarily large, paired datasets using models which were trained on few or no paired samples.

In this work, we demonstrate that the benefits of training on large, paired datasets can be replicated using synthetically generated paired data for privileged self-supervised training. We show that high quality, biologically informed self-supervised models can be trained that benefit from privileged information, but only require a single data source for training and inference by using that input to generate a synthetic privileged pair. The key contributions of this work are:

- We demonstrate that synthetically generated data can be used as privileged information in place of authentic source data, including where there is little or no real paired data, to improve model performance;
- We show that these models are more robust to distribution shift than un-privileged models, and encode meaningful biological features;
- We demonstrate that models trained with synthetic data closely mimic the representations learned from real data, and learn extra information compared to training with a single modality;
- We demonstrate the extra benefit from distilling synthetically generated data when the real dataset size is limited.

2   Background and Related Work

Self-Supervised Learning (SSL) is a midpoint between supervised and unsupervised, where inherent structural properties of data are leveraged to to create inductional biases which enable the model to learn without labels. The goal of this is to produce a vector representation of the image which has a much lower dimension than the original image but contains the useful semantics of the image needed for downstream applications. Recent work has extensively applied SSL in a wide variety of domains and tasks [7], and state-of-the-art SSL models have been shown to outperform both supervised and unsupervised models in many settings.

Privileged Information is any information that is available during training but not during inference. In a medical imaging context, this is true of an enormous corpus of scientific data, which has already been generated in one condition, but which would be impractical or expensive to generate again for future studies or to create a larger dataset. We are primarily motivated by histopathology, where vast quantities of brightfield H&E stain datasets are available, but comparatively few for immunohistochemistry (IHC) stains and fewer again for state-of-the-art data such as spatial -omics, multiplex IHC, super-resolution microscopy or expert pixel-level annotations.

As Siamese networks will generally learn to encode only features which are shared between branches, if the inputs are different then only semantic content shared between inputs is learned. This means that useful features can be neglected if they are not shared. For example, it was shown that pairing thresholded immunofluorescence (IF) images with H&E images significantly decreased model downstream performance on the H&E images because the model neglected many details, such as fine-grained tissue structure, in these images that were not detectable in the IF, which were mostly black [1].

TriDeNT [1] was developed to address this issue by only neglecting features that are not present in the primary image. This means that features may still be neglected by the encoder of the privileged input, but only the encoder of the primary input is of interest for downstream tasks. TriDeNT works by utilising an additional primary input branch to generate a third representation, meaning that there are three concurrent joint-embedding latent spaces. The model must then balance retaining features that are present in both primary branches with those shared by the privileged information, and was shown to retain both feature sets, significantly improving downstream performance.

Image Generation methods have been extensively applied in medical imaging, both generating realistic images from a random seed [8], or virtually translating between modalities [9] or stains [10,11]. Some image methods require paired data, such as pix2pix [12], but others, such as cycleGAN [6], can find mappings between unpaired datasets. This is especially useful in the context of histology, as the vast majority of additional stains are not from the same slice as their paired H&E stains, and often are from different sections of the tissue block. This means it is hard to determine which parts of the images should be aligned when comparing different stains, as their morphologies can be very different. It

has been shown that CycleGANs can be used to generate realistic IHC stains directly from H&E stains without requiring corresponding input/target pairs [13,10], and we demonstrate here that these synthetically generated stains can be used to improve the performance of self-supervised models using TriDeNT. These stains usually have the advantage of being perfectly spatially registered and aligned.

It has been shown that the representations from generative models perform poorly when evaluated on downstream tasks [14], with even support vector machines outperforming state-of-the-art generative models [15]. This is due to the models being highly optimised for the specific task of image generation, and therefore having representations which contain information, such as precise spatial locations of features, which are not useful for downstream tasks such as tissue type classification. It is therefore essential for representation learning to distil the useful information from these models' representations without over-constraining the model to replicate all of the features present, which reduces the quality of representations.

3   Methodology and Datasets

In all tasks we follow the implementation details of [1]. We compare TriDeNT to Siamese baselines, with and without privileged information, and also to supervised learning. Encoders are ResNet-50 [17] trained for 100 epochs with a warmup-cosine learning rate, with maximum value 10^{-4}. Three-layer dense projection heads of width 8192 with batch normalisation were used, and classifier heads were single dense layers with softmax activations. For full details of datasets used, see Table S1. We employ three evaluation tasks: tissue classification on the NCT dataset [18], metastasis detection on the Camelyon17 dataset [19], and malignancy detection on the PanNuke dataset [20]. We use existing image generation models trained on PanNuke, SHIFT [4] and a pix2pix model and a cycleGAN both trained on the IHC dataset detailed in [10]. The pix2pix model uses paired image-to-image translation, while the cycleGAN uses unpaired images, allowing the necessity for paired data in model training to be assessed.

4   Results

4.1   Synthetically Generated Segmentations as Privileged Data Improves Downstream Classification Performance

In [14], it was shown that standard Siamese unprivileged models failed to learn sub-nuclear features from histopathology images. This was addressed by pairing these images with machine-generated nuclear segmentations, which prompted the model to learn features related specifically to the nuclei. This was then further explored in [1], where it was shown that training models using TriDeNT on the Pannuke dataset [20] with the segmentations as privileged information greatly improves their downstream task performance. Here we use the setting of
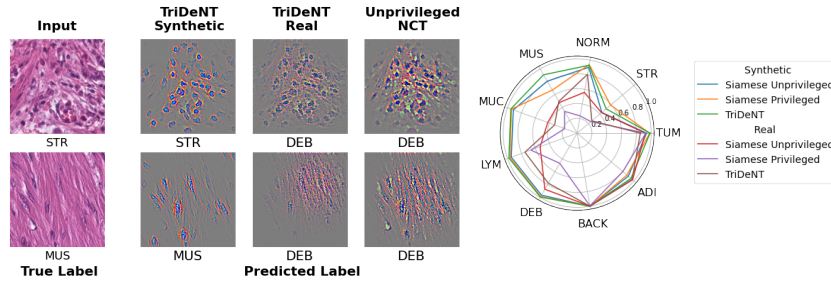
(a)

(b)

Loss	Method	Privileged	PanNuke		IHC (cycleGAN)		IHC (pix2pix)		SHIFT	
			NCT	Camelyon	NCT	Camelyon	NCT	Camelyon	NCT	Camelyon
VICReg	TriDeNT	✓	**0.9359**	**0.9149**	**0.9252**	0.9263	0.9184	0.8960	**0.9383**	0.6294
		✓	0.9044	0.8952	0.9234	**0.9462**	**0.9250**	**0.8994**	0.9236	0.5975
	Siamese	✗ NCT	0.8937	0.7056	0.8937	0.7056	0.8937	0.7056	0.8937	**0.7056**
		✗ Real	0.7301†	0.4636	0.8434	0.5360	0.8434	0.5360	0.8008	0.6387
InfoNCE	TriDeNT	✓	**0.9287**	0.8719	**0.9331**	0.9033	**0.9302**	0.8992	0.9115	0.6430
		✓	0.9103	0.8889	0.9260	**0.9356**	0.9259	**0.9224**	**0.9299**	0.6007
	Siamese	✗ NCT	0.9047	0.6205	0.9047	0.6205	0.9047	0.6205	0.9047	0.6205
		✗ Real	0.7199†	0.5852	0.8623	0.5600	0.8623	0.5600	0.7921	**0.7056**
CrossEntropy	Supervised	✗	0.9245	0.8440	0.9245	0.8440	0.9245	0.8440	0.9245	0.8440

(c)

Loss	Method	Privileged	PanNuke				IHC (cycleGAN)				SHIFT			
			NCT		Camelyon		NCT		Camelyon		NCT		Camelyon	
			Synthetic	Real	Synthetic	Real	Synthetic	Real	Synthetic	Real	Synthetic	Real	Synthetic	Real
VICReg	TriDeNT	✓	**0.9359**	0.7337†	**0.9149**	0.5306	**0.9252**	0.8650	**0.9263**	0.6266	**0.9383**	0.7624	**0.6294**	0.5339
	Siamese	✓	**0.9044**	0.6000†	**0.8952**	0.4831	**0.9234**	0.8256	**0.9462**	0.8614	**0.9236**	0.7439	**0.5975**	0.5120
InfoNCE	TriDeNT	✓	**0.9287**	0.7530†	**0.8719**	0.6908	**0.9331**	0.8722	**0.9033**	0.8304	**0.9115**	0.7998	**0.6430**	0.5653
	Siamese	✓	**0.9103**	0.5289†	**0.8889**	0.7134	**0.9260**	0.8370	**0.9356**	0.8423	**0.9299**	0.7464	0.6007	**0.6601**

(d)

Fig. 2: (a) Guided-GradCAM [16] maps for samples from the NCT dataset. Predicted labels for each classifier are shown, as is the true label. Labels are adipose (ADI), background (BACK), debris (DEB), lymphocytes (LYM), mucus (MUC), muscle (MUS), normal mucosa (NORM), stroma (STR), and tumour (TUM). See Figure S1 for examples of all classes; (b) Breakdown of performance on the NCT evaluation task by class; (c) Results for models trained on the NCT dataset with synthetically generated privileged information as pairs for the NCT patches. Values marked with "†" are from [1]; (d) Classification performance on NCT and Camelyon of models trained on the PanNuke/IHC/SHIFT datasets compared to the synthetically generated images paired with the NCT dataset.

[14], generating segmentation maps using HoVer-Net [21], and study the effect of using this as privileged information for TriDeNT. Note that PanNuke contains train/validation/test folds of approximately 2000 images each, meaning the SSL training dataset size is substantially increased by generating synthetically generated data from the NCT dataset.

The results in Figure 2 demonstrate that distilling information from the synthetically generated nuclear segmentations greatly improves the performance of the model, and synthetically generated privileged information is found to be effectively distilled into the primary models. Notably, unprivileged learning on PanNuke or the NCT dataset performed poorly at distinguishing stroma from muscle, which requires subtle biological features to be learned. Models with synthetically generated privileged information performed considerably better. Encoders trained with privileged information retain features of both the background/connective tissue (using information from only the H&E patches) and the nuclei (using information from both the H&E patches and the segmentation masks), enabling them to make accurate predictions across different tissue types. Figures 2a and S1 demonstrate that the TriDeNT model with synthetically generated privileged information learns to accurately identify the morphology of the nuclei in the images, making it easier to identify the phenotypic change between muscle and stroma. It was previously observed that unprivileged Siamese models tend to focus on the background, rather than the potentially more informative features of the nuclei [1]. Privileged training with segmentation masks produced models which focused more on nuclei, this is equally true of synthetically generated privileged information. We find that even providing only the segmented nuclei as privileged information achieves improvements, although to a lesser extent than also providing the type annotations, as shown in Table S2.

We note that there is a marked performance improvement for models with synthetically generated privileged information on the Camelyon17 task, which assesses the ability of the model to both detect metastasis and transfer to out-of-distribution datasets. Performance on this task is a good indicator of whether meaningful, generalisable biological features have been learned, as opposed to features specific to the train set or simply the contours of an image. It is known that privileged information causes in models to learn representations which are less prone to overfitting on small changes in irrelevant features such as colour, brightness, or blur, which can easily arise between datasets [1]. Instead, models must focus on more useful semantic features shared between primary and privileged branches, such as the nuclei in Figures 2a and S1, which naturally gives rise to representations which transfer well to out-of-distribution data. We find that this is equally true for synthetic privileged information.

We also find that training on synthetically generated training sets offers significant improvements compared to using privileged information from the real source dataset, as shown in Figure 2d. This is due to the synthetic dataset being either larger (compared to PanNuke and SHIFT) or being more curated to encompass a distribution of samples (compared to IHC). This demonstrates that not only is using synthetic privileged information a good method to enhance performance on tasks related to the dataset the synthetic privileged information is generated *for*, but it can also offer better results on tasks relating to the dataset the synthetic privileged information was generated *from*.

As above, we see significant performance improvements for both pix2pix and cycleGAN models, showing that no real paired data is needed to achieve good
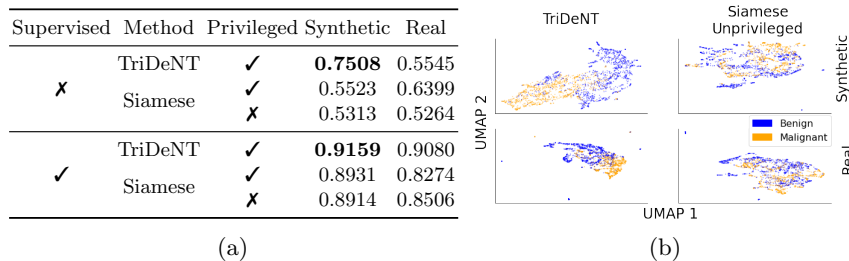
Supervised	Method	Privileged	Synthetic	Real
✗	TriDeNT	✓	**0.7508**	0.5545
	Siamese	✓	0.5523	0.6399
		✗	0.5313	0.5264
✓	TriDeNT	✓	**0.9159**	0.9080
	Siamese	✓	0.8931	0.8274
		✗	0.8914	0.8506

(a)   (b)

Fig. 3: (a) Evaluation of representations from models trained on the real PanNuke dataset and the synthetically generated nuclear segmentations paired with the NCT dataset. (b) UMAP projections, coloured by PanNuke label.

results. It is notable that the synthetic SHIFT dataset improves performance on the NCT task, although this does not generalise well to the Camelyon evaluation task. The SSL performance is improved in spite of the generator making many verifiable errors, such as predicting fluorescence in patches labelled "background" or "debris", which makes little biological sense. This demonstrates that the addition of any useful privileged information can enhance performance, and synthetically generated data does not need to be perfect, with even models trained on very small datasets being suitable.

4.2 Training With Synthetically Generated Privileged Data Outperforms Training on the Real Source Dataset

Another key use case for synthetically generated privileged information is where only a small amount of data is available. In Figure 3a, we evaluate models trained on the small PanNuke dataset by performing unsupervised k-means clustering with 2 clusters to determine binary labels. Trained with synthetically generated privileged information, TriDeNT was able to accurately make completely unsupervised classifications. For all other methods, performance is little better than random, implying that neither the PanNuke train set nor unprivileged training on the larger dataset is sufficient for good performance. This can be intuitively seen from Figure 3b, where TriDeNT differentiates between the two classes more than unprivileged methods, with qualitative similarities between the projections of real and synthetic models. This shows that this approach could be used to better analyse small datasets from rare diseases or underrepresented populations, which may not have sufficient labels for classifier training, either for basic research or in a clinical setting. Supervised evaluation produces similar results, with the TriDeNT model trained on NCT with synthetically generated PanNuke patches outperforming training directly on the real PanNuke dataset.

It is notable that models performed well which were trained with synthetically generated privileged information from small datasets such as PanNuke and SHIFT, which have generator train sets of 2205 and 665 images respectively. Figure 2d demonstrates that the reduction in error from training on the larger

synthetic dataset compared to the real smaller dataset is more than 400% and 200% for PanNuke and SHIFT respectively.

5   Discussion and Conclusions

5.1   Synthetically Generated Privileged Data Improves Model Performance in Both In- and Out-of-Distribution Settings

We have established that the performance of single-input models is significantly improved by training with synthetically generated paired data. This can be considered as a form of implicit multi-objective learning, as the image generation model has been optimised to perform a different task to the self-supervised pretext task, and this provides the commensurate benefits. Despite this, the required computational power at any point is no different to standard Siamese/TriDeNT learning, as the images are generated prior to the self-supervised training. In many cases, these generative models are already available and could be used to create images without being trained. This means that these methods can be accessible to clinicians and researchers without large-scale resources, and enable the application of self-supervised privileged training to datasets which are otherwise intractable due to a lack of paired data.

5.2   Clinical Impact and Future Work

This work was primarily motivated by two use cases: unpaired datasets, and small/rare datasets. Many state-of-the-art datasets do not have accompanying routine data, and this approach allows insights from those data to be integrated into routine data. Furthermore, many datasets are very limited in size or scope, particularly for rare diseases or those affecting underrepresented populations. This approach could enable retrospective reanalysis of routine data from rare samples using insights from current state-of-the-art data, meaning routine data from both rare and common cases can be fully utilised. This is especially prescient, given that TriDeNT was shown to significantly outperform comparable methods for few-shot applications [1].

We have limited the scope of this study to the case of a single paired input, but further improvements could be made with the addition of multiple paired datasets. This could be easily achieved by applying several generative models to the same large primary dataset. We have also not considered training with the real data shuffled into the synthetically generated dataset, which could further improve performance.

Acknowledgements

References

1. Lucas Farndale, Robert Insall, and Ke Yuan. Trident: Triple deep network training for privileged knowledge distillation in histopathology. *arXiv preprint arXiv:2312.02111*, 2023.
2. Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190, 2023.
3. Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
4. Erik A Burlingame, Mary McDonnell, Geoffrey F Schau, Guillaume Thibault, Christian Lanciault, Terry Morgan, Brett E Johnson, Christopher Corless, Joe W Gray, and Young Hwan Chang. Shift: speedy histological-to-immunofluorescent translation of a tumor signature enabled by deep learning. *Scientific reports*, 10(1):1–14, 2020.
5. Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jegou, and Edouard Grave. Are large-scale datasets necessary for self-supervised pretraining? *arXiv preprint arXiv:2112.10740*, 2021.
6. Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
7. Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):4037–4058, 2020.
8. Amirhossein Kazerouni, Ehsan Khodapanah Aghdam, Moein Heidari, Reza Azad, Mohsen Fayyaz, Ilker Hacihaliloglu, and Dorit Merhof. Diffusion models in medical imaging: A comprehensive survey. *Medical Image Analysis*, page 102846, 2023.
9. Qing Lyu and Ge Wang. Conversion between ct and mri images using diffusion and score-matching models. *arXiv preprint arXiv:2209.12104*, 2022.
10. Christopher D Walsh, Joanne Edwards, and Robert H Insall. Ensuring accurate stain reproduction in deep generative networks for virtual immunohistochemistry. *arXiv preprint arXiv:2204.06849*, 2022.
11. Kevin de Haan, Yijie Zhang, Jonathan E Zuckerman, Tairan Liu, Anthony E Sisk, Miguel FP Diaz, Kuang-Yu Jen, Alexander Nobori, Sofia Liou, Sarah Zhang, et al. Deep learning-based transformation of h&e stained tissues into special stains. *Nature communications*, 12(1):4884, 2021.
12. Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
13. Zhaoyang Xu, Xingru Huang, Carlos Fernández Moro, Béla Bozóky, and Qianni Zhang. Gan-based virtual re-staining: a promising solution for whole slide image analysis. *arXiv preprint arXiv:1901.04059*, 2019.
14. Lucas Farndale, Robert Insall, and Ke Yuan. More from less: Self-supervised knowledge distillation for information-sparse histopathology data. *arXiv preprint arXiv:2303.10656*, 2023.
15. Adalberto Claudio Quiros, Nicolas Coudray, Anna Yeaton, Wisuwat Sunhem, Roderick Murray-Smith, Aristotelis Tsirigos, and Ke Yuan. Adversarial learning of cancer tissue representations. In *Medical Image Computing and Computer Assisted*

Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VIII 24, pages 602–612. Springer, 2021.

16. Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.

17. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

18. Jakob Nikolas Kather, Niels Halama, and Alexander Marx. 100,000 histological images of human colorectal cancer and healthy tissue, April 2018.

19. Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermsen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 2018.

20. Jevgenij Gamper, Navid Alemi Koohbanani, Ksenija Benes, Ali Khuram, and Nasir Rajpoot. Pannuke: an open pan-cancer histology dataset for nuclei instance segmentation and classification. In *European Congress on Digital Pathology*, pages 11–19. Springer, 2019.

21. Simon Graham, Quoc Dang Vu, Shan E Ahmed Raza, Ayesha Azam, Yee Wah Tsang, Jin Tae Kwak, and Nasir Rajpoot. Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Medical Image Analysis*, 58:101563, 2019.

22. Zhi Huang, Federico Bianchi, Mert Yuksekgonul, Thomas J Montine, and James Zou. A visual–language foundation model for pathology image analysis using medical twitter. *Nature Medicine*, pages 1–10, 2023.

Supplementary Material

Table S1: Datasets

Name	Tissue Type	Train/Test Split	Training	Evaluation	Purpose	Task
NCT [18]	Colorectal	100,000/7,177	✓	✓	Cancer Tissue Analysis	Tissue Type Classification

Note: Nine tissue types: adipose (ADI), background (BACK), debris (DEB), lymphocytes (LYM), mucus (MUC), smooth muscle (MUS), normal colon mucosa (NORM), cancer-associated stroma (STR), colorectal adenocarcinoma epithelium (TUM)

Name	Tissue Type	Train/Test Split	Training	Evaluation	Purpose	Task
Camelyon17 [19]	Lymph Node	179,394/146,722	✗	✓	Out-of Distribution Metastasis Detection	Out-of-Distribution Binary Classification

Note: Train set from 3 centers, validation set from center 4, and test set from center 5

Name	Tissue Type	Train/Test Split	Training	Evaluation	Purpose	Task
PanNuke [20]	Pan-Cancer	2,205/4,373	✓	✓	Nuclear Segmentation	Neoplastic Cell Detection

Note: As in [22,1], patches are labelled as benign if no neoplastic cells are present, and malignant if the patch contains at least 10 cells and 30% of the cells are neoplastic.

Name	Tissue Type	Train/Test Split	Training	Evaluation	Purpose	Task
Synthetic IHC [10]	Breast	93,952/-	✓	✗	Virtual Restaining Brightfield AE1/AE3	-

Note: Originally 23,488 448 × 448px patches, which were split into 4 smaller patches

Name	Tissue Type	Train/Test Split	Training	Evaluation	Purpose	Task
SHIFT [4]	Prostate	665/11,593	✓	✗	Virtual Restaining IF pan-CK	-

Note: All 12,258 patches used in SSL training. Generator trained on only 665. Only 4 patients' data in dataset.

Table S2: Results for models trained with synthetically generated nuclear segmentations as privileged information for the NCT dataset. Unlike in the results in the main text, these segmentations simply mask the background, leaving the nuclei. There are no type annotations. Evaluation is performed on the NCT and Camelyon datasets.

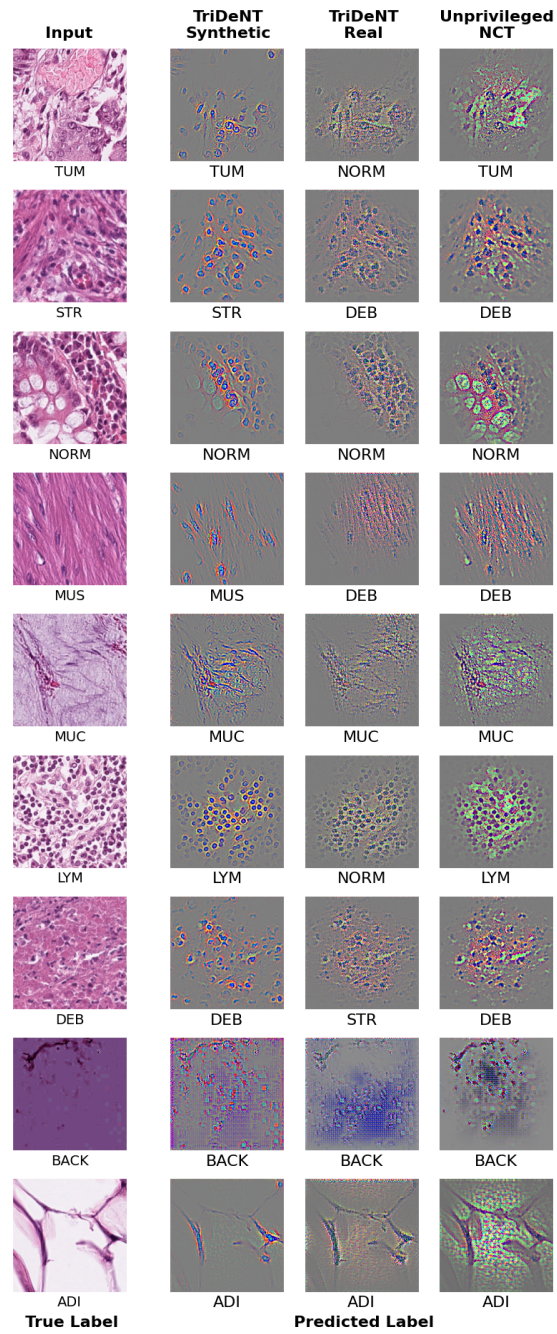| Loss	Method	Privileged	NCT	Camelyon
VICReg	TriDeNT	✓	**0.9280**	**0.8330**
	Siamese	✓	0.7778	0.7871
		✗	0.8937	0.7056
InfoNCE	TriDeNT	✓	**0.9334**	0.6859
	Siamese	✓	0.9058	**0.8878**
		✗	0.9047	0.6205

Fig. S1: Representative samples from each NCT tissue class with Guided-GradCAM activation maps. True and predicted labels are shown for each classifier.