Claudio Quiros, A., Coudray, N., Yeaton, A., Sunhem, W., Murray-Smith, R., Tsirigos, A. and Yuan, K. (2021) Adversarial Learning of Cancer Tissue Representations. In: 24th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2021), 27 Sept-1 Oct 2021, pp. 602-612. ISBN 9783030872366 (doi:[10.1007/978-3-030-87237-3_58](https://doi.org/10.1007/978-3-030-87237-3_58)).

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/250847/>

Deposited on: 03 September 2021

Adversarial learning of cancer tissue representations

Adalberto Claudio Quiros¹ ✉, Nicolas Coudray², Anna Yeaton², Wisuwat Sunhem¹, Roderick Murray-Smith¹, Aristotelis Tsirigos², and Ke Yuan¹ ✉

¹ University of Glasgow School of Computing Science, Scotland, UK

² New York University School of Medicine, New York, NY, USA

a.claudio-quiros.1@research.gla.ac.uk, nicolas.coudray@nyulangone.org,
anna.yeaton@nyulangone.org, w.sunshem.1@research.gla.ac.uk,
roderick.murray-smith@glasgow.ac.uk
aristotelis.tsirigos@nyulangone.org, ke.yuan@glasgow.ac.uk

Abstract. Deep learning based analysis of histopathology images shows promise in advancing the understanding of tumor progression, tumor micro-environment, and their underpinning biological processes. So far, these approaches have focused on extracting information associated with annotations. In this work, we ask how much information can be learned from the tissue architecture itself.

We present an adversarial learning model to extract feature representations of cancer tissue, without the need for manual annotations. We show that these representations are able to identify a variety of morphological characteristics across three cancer types: Breast, colon, and lung. This is supported by 1) the separation of morphologic characteristics in the latent space; 2) the ability to classify tissue type with logistic regression using latent representations, with an AUC of 0.97 and 85% accuracy, comparable to supervised deep models; 3) the ability to predict the presence of tumor in Whole Slide Images (WSIs) using multiple instance learning (MIL), achieving an AUC of 0.98 and 94% accuracy.

Our results show that our model captures distinct phenotypic characteristics of real tissue samples, paving the way for further understanding of tumor progression and tumor micro-environment, and ultimately refining histopathological classification for diagnosis and treatment. ³

Keywords: Generative Adversarial Networks · Histology

1 Introduction

Histological images, such as hematoxylin and eosin (H&E) stained tissue microarrays (TMAs) and whole slide images (WSIs), are an imaging technology routinely used in clinical practice, that relay information about tumor progression and tumor microenvironment to pathologists. Recently, there has been advances in relating tissue phenotypes from histological images to genomic mutations [2–4, 10], molecular subtypes [20, 21], and prognosis [3, 4]. While these

³ The code and pretrained models are available at: <https://github.com/AdalbertoCq/Adversarial-learning-of-cancer-tissue-representations>

previous works have contributed greatly to the field, and have highlighted the richness of information in H&E slides, they are limited by the frequently unavailable labels required to train supervised methods.

Unsupervised learning offers the opportunity to build phenotype representations based on tissue architectures and cellular attributes, without expensive labels or representations that are only correlated with a selected predicted outcome as discriminative models.

Within unsupervised models, Generative Adversarial Networks (GANs) have been widely used in digital pathology, from nuclei segmentation [14], stain transformation and normalization [18, 23], to high-quality tissue samples [12]. In addition, there has been some initial work on building representations of cells [5] or larger tissue patches [16].

In particular, PathologyGAN [16] offers a Generative Adversarial Network (GAN) that captures tissue features and uses these characteristics to give structure to its latent space (e.g. colour, texture, spatial features of cancer and normal cells, or tissue type). The generator has representation learning properties where distinct regions of the latent space are directly related to tissue characteristics. However, it does not offer the ability to project real tissue images to its latent space, only offering latent representations for synthesized tissue. Such representations are crucial if we want to relate morphological and cellular features of tissue with genetic, molecular, or survival patient information.

Contributions We propose an adversarial learning model to extract feature representations of cancer tissue, we show that our tissue representations are built on morphological and cellular characteristics found on real tissue images. We present a Generative Adversarial Network with representation learning properties that includes an encoder, allowing us to project real tissue onto the model’s latent space. In addition, we show evidence that these representations capture meaningful information related to tissue, illustrating their applicability in three different scenarios:

- 1) We visualize the representations of real tissue and show how distinct regions of the latent space enfold different characteristics, and reconstructed tissue from latent representations follow the same characteristics as the original real image; 2) We quantify how informative these representation are by training a linear classifier to predict tissue types; 3) We demonstrate the suitability of the tissue representations on a multiple instance learning (MIL) task, predicting presence of tumor on WSIs.

2 Method

We build upon PathologyGAN [16], and introduce an encoder E that maps back generated tissue to the GAN’s latent space. Effectively learning to interpret tissue morphologies and acting as the inverse of the generator G . After training, the encoder can be used independently to map real tissue samples to their rep-

representations in the latent space. Figure 1 captures the high level architecture of our model.

The loss functions for the discriminator L_{Dis} and the generator L_{Gen} are defined as the Relativistic Average Discriminator [7], where the discriminator's goal is to estimate the probability of the real data being more realistic than the fake (Equations 1 and 2). P_{data} is the distribution of real data, $G(w)$ is the distribution of synthetic data produced by the Generator G , $C(x)$ is the non-transformed discriminator output or critic, W the transformed latent space with representation learning properties, and $P_z = \mathcal{N}(0, I)$ the original latent space:

$$L_{Dis} = -\mathbb{E}_{x_r \sim P_{data}} \left[\log \left(\tilde{D}(x_r) \right) \right] - \mathbb{E}_{x_f \sim G(w)} \left[\log \left(1 - \tilde{D}(x_f) \right) \right], \quad (1)$$

$$L_{Gen} = -\mathbb{E}_{x_f \sim G(w)} \left[\log \left(\tilde{D}(x_f) \right) \right] - \mathbb{E}_{x_r \sim P_{data}} \left[\log \left(1 - \tilde{D}(x_r) \right) \right], \quad (2)$$

$$\tilde{D}(x_r) = \text{sigmoid} \left(C(x_r) - \mathbb{E}_{x_f \sim G(w)} C(x_f) \right),$$

$$\tilde{D}(x_f) = \text{sigmoid} \left(C(x_f) - \mathbb{E}_{x_r \sim P_{data}} C(x_r) \right),$$

$$w = M(z), \quad z \sim P_z.$$

We use the mean square error between latent vectors w and their reconstruction from generated images $w' = E(G(w))$ as the encoder loss function, L_{Enc} (Equation 3):

$$L_{Enc} = \mathbb{E}_{z \sim P_z} \left[\frac{1}{n} \sum_{i=1}^n (w_i - w'_i)^2 \right] \text{ where } w' = E(G(w)), \quad w = M(z). \quad (3)$$

Although the encoder E is simultaneously trained with the GAN model, we can separate the model training into two parts: The mapping network M , generator G , and discriminator D that are trained as a GAN, and the encoder E , which is trained to project back the generated cancer tissue images onto the latent space. In practice, the encoder E learns with the Generator G . We trained our encoder based on the assumption that the generator is successful in reproducing cancer tissue. Therefore the encoder will learn to project real tissue images if it is able to do so with generated ones. Based on this logic, we use only generated images to train the encoder.

Fig. 1. High level architecture of our GAN model.

Additionally, the encoder is only updated when the generator is not trained with style mixing regularization [8]. Style mixing regularization is a technique that promotes disentanglement where the generator G generates an image from two different latent vectors w_1 and w_2 , these vector are feed at different layers of the generator G . However, it becomes impractical to train the encoder in these steps because these images have no clear assignation in the latent space W . Therefore, The style mixing regularization is only preformed 50% of times in the generator training, so our encoder is updated every two steps per the generator.

Datasets We trained our model with three different cancer types: Breast, colon, and lung. We provide a detailed description on these datasets and how they were built in the Appendix A.

The breast H&E cancer dataset is composed by the Netherlands Cancer Institute (NKI, Netherlands) and Vancouver General Hospital (VGH, Canada) cohorts [1] with 248 and 328 patients, each patient with associated Tissue Micro-Arrays (TMAs). Each TMA is divided into 224×224 tiles and labeled subject to density of cancer cells in the tile using CRImage [22], we use 9 different classes with class 8 accounting for tiles with the largest count of cancer cells. This dataset is composed by a training set of 249K tiles and 460 patients, and a test set of 13K tiles and 116 patients, with no overlapping patients between sets.

The colorectal H&E cancer dataset from National Center for Tumor diseases (NCT, Germany) [9] provides tissue images of 224×224 resolution with an associated type of tissue label: Adipose, background, debris, lymphocytes, mucus, smooth muscle, normal colon mucosa, cancer-associated stroma, and colorectal adenocarcinoma epithelium (tumor). The dataset is divided into a training set of 100K tissue tiles and 86 patients, and a test set of 7K tissue tiles and 50 patients, there is no overlapping patients between train and test sets. Finally, in order to compare our model to supervised methods on tissue type classification, we combine the classes stroma and smooth muscle into a class 'simple stroma' as in reference publications [11, 17].

The lung H&E cancer dataset contains samples with adenocarcinoma (LUAD), squamous cell carcinoma (LUSC), and normal tissue, composed by 1807 Whole Slide Images (WSIs) of 1184 patients from the Cancer Genome Atlas (TCGA). Each WSI is labeled as tumor and non-tumor depending on the presence of lung cancer in the tissue and divided into 224×224 tiles. In addition, we split the dataset into a training set of 916K tissue tiles and 666 patients, and a test set of 569K tissue tiles and 518 patients, with no overlapping patients between both sets.

3 Results

Since our tissue representations are built in an unsupervised manner and from tissue features alone, we studied how meaningful real tissue representations are. We show results quantifying the tissue morphology and cellular information contained in the representations, and provide examples of how they can be exploited:

Latent space visualization and tissue reconstruction, tissue type classification, and tumor prediction in a multiple instance learning setting.

3.1 Visualizations of tissue representations and reconstructions

We first aim to test how interpretable the latent representations of real tissue are, and if they capture enough information to recreate tissue with the same characteristics. We do so by visualizing the latent representations along with the corresponding real images and by reconstructing images from real tissue projections.

We used the breast cancer and colorectal cancer datasets for these results. As covered in Section 2, the breast tissue patches have an associated label with the density of cancer cells in the image and the colorectal tissue patches have an associated label with the corresponding tissue type. In both cases, we used the training set to train our GAN model and later used the encoder E to project images from test sets, obtaining representations of real tissue samples.

In Figure 2 we used UMAP [15] to reduce the dimensionality of latent representations from 200 to 2 and plot each latent vector with its associated real tissue image and label. In the breast cancer case, tissue with the highest density of cancer cells is distributed outwards while no presence of them concentrates inwards. Colorectal cancer samples are located in different regions of the space depending of their tissue types. In both examples, tissue characteristics of real images, whether if it is tissue type or cancer cell density, determines the location of the representation in the latent space.

In addition, we wanted to compare real images with the synthetic images reconstructed by the generator at the specific position in the latent space where each real image was projected, $X_{recon} = G(E(X_{real}))$. Figure 3 shows real images and their associated reconstructions, we used latent representations of real images (a) to synthesize their reconstructions (b) with the GAN’s generator. We provide examples of the colorectal cancer tissue types (1-5) and different cancer cell densities in breast cancer tissue (6-10). We can see that the reconstructions follow the characteristics of the real tissue images.

We conclude that tissue representations are not only interpretable but also hold the relevant information needed to synthesize tissue with the same characteristics. We provide further samples of latent space visualizations in the Appendix B

3.2 Tissue type classification over latent representations

In this task, we aim to quantify how informative latent representation are, verifying that the model is able to learn tissue patterns that define tissue type. We train a linear classifier to predict tissue type over latent representations.

We used the colorectal cancer dataset with each tissue sample and its associated tissue type. We trained our GAN with the training set and again used the Encoder E to project the training and test sets onto the GAN’s latent space, obtaining representations for each tissue image. Consecutively, we used the latent

representations for the training set to train a logistic regression (one-versus-rest) and evaluate performance using test set projections.

Table 1 shows the performance of the logistic regression, reaching an average AUC (one-vs-rest) of 0.976 and multi-class accuracy of 85.43% just from latent representations. Table 2 provides a comparison between other existing supervised methods, such as Raczkowski et al. [17] where a Bayesian deep neural network

Fig. 2. Uniform Manifold Approximation and Projection (UMAP) vectors of PathologyGAN’s latent representations, we present breast cancer tissue from NKI and VGH (left image) and colorectal cancer tissue from NCT (right image). Breast cancer tissue images are labeled using cancer cell counts, class 8 accounting for the largest cell number. Colorectal cancer tissue images are labeled based on their tissue type. In both cases, we observe that real tissue projected to the latent space retain cellular information and tissue morphology.

Fig. 3. Real tissue images and their reconstructions. We take real tissue images and map them to the latent space with our encoder, then we use the generator with the latent vector representations to generate the image reconstructions. (a) correspond to the real tissue images X_{real} and (b) to the reconstructions $X_{recon} = G(E(X_{real}))$, the images are paired in columns. We present samples of colorectal cancer tissue from NCT (1-5) and breast cancer tissue from VGH and NKI (5-10). We show different examples of tumor(1,9,10), stromal (2,8), normal mucosa (3), muscle (3), and lymphocytes (6,7), the reconstructions follow the real image attributes.

Total AUC: 0.976							
Tumor	Simple	Stroma	Mucosa	Mucus	Lymph.	Debris	Back. Adipose
0.974	0.929		0.964	0.997	0.994	0.959	1.0 0.998
Total Accuracy: 85.43%							
Tumor	Simple	Stroma	Mucosa	Mucus	Lymph.	Debris	Back. Adipose
89%	71%		68%	91%	83%	63%	100% 96%

Table 1. Logistic regression Accuracy and AUC on tissue type classification. We used the colorectal cancer tissue images from NCT to train a classifier over latent representations. A logistic regression (one-vs-rest) is able to reliably find information in tissue representations to predict tissue types.

Model	AUC	Accuracy
Ours	0.976	85.43%
Bayesian DNN [17]	0.995	99.2%
RBF-SVM [11]	0.976	87.4%

Table 2. Performance comparison on tissue type classification between existing methods and our tissue latent representations. These results reflect baseline performance of supervised deep learning and non-deep learning approaches. Our representations on logistic regression without any transformation or projection, and the fact that they are comparable with supervised performance, demonstrate the applicability and information hold in them.

is trained for the purpose of tissue classification achieving an AUC of 0.992 and 92.44% accuracy, in addition, Kather et al. [11] provides another performance reference using an RNF-SVM with an AUC of 0.995 and 92.44% accuracy. We provide further details on ROC and confusion matrices in the Appendix D

Phenotype representations from our model provide robust information such that a linear classifier is able to predict tissue types with high accuracy. Despite a small loss of performance when compared to the best supervised method, our fully unsupervised representations could be an effective solution when extensive and precise manual annotations cannot be obtained.

3.3 Multiple Instance Learning on latent representations

Finally, we tested the reliability of the tissue phenotype representations in a weakly supervised setting. We used tissue representations in a multiple instance learning (MIL) task for tumor presence prediction in Whole Slide Images (WSIs), where each WSI has an average, minimum, and maximum of 974, 175, and 14K, tiles, respectively.

We used the lung cancer dataset where each WSI has an associated label tumor or normal subject to presence of tumor tissue. We divided each WSI into 224×224 patches to train our GAN model and obtain tissue representations, and we later used all tissue representations of the WSI.

In the case of the MIL problem, we have a bag of instances $X = \{x_1, \dots, x_n\}$ with an individual label $Y \in \{0, 1\}$. We further assume that each instance of the

bag has an associated label $y_i \in \{0, 1\}$ to which we have no access:

$$Y = \begin{cases} 0, & \text{iff } \sum_k y_k = 0 \\ 1, & \text{otherwise} \end{cases}$$

In our case, we have tissue tiles that we translated into tissue representations $W = \{w_1, \dots, w_n\}$, and use them to determine the presence of lung tumor $Y = 1$ in the WSI. We used the attention-based deep MIL [6] as it assigns a weight to each instance of the bag x_k , allowing us to measure which representations are relevant for the prediction.

The attention-based deep MIL over latent representations achieves an AUC of 0.98 and accuracy of 94%. These results are comparable to Coudray et al. [2], where an Inception-V3 network was tested on the same dataset and purpose achieving an AUC of 0.993 and accuracy of 97.5%.

Fig. 4. Uniform Manifold Approximation and Projection (UMAP) vectors of lung cancer tissue representations. We labeled each patch of the WSI with the corresponding label subject to presence of tumor in the WSI, and highlight images and representations where the attention-based deep MIL focuses to predict the outcome. We can see that the MIL framework emphasizes on regions where there is only the presence of tumor patches, relying on the information contained in the latent representations.

Figure 4 shows UMAP reductions of tissue patches representations for the test set, we labeled each patch representation with the WSI label, tumor or normal depending on the presence of tumor. In addition we highlighted images and representations of the top 0.1% most weighted representations for the tumor outcome prediction. We can see that the MIL model focuses on regions of the latent space where there is no presence of normal tissue.

Given the accurate prediction of lung tumor presence in the WSIs and the focus of the MIL framework on representations solely present in tumor tissue, we conclude that the phenotype representations are reliable enough to capture tissue characteristics such as tumor tissue at scale.

4 Conclusion

We presented an adversarial learning model that builds phenotype representations of cancer tissue. Our model overcomes previous limitations [16] by introducing an Encoder, enabling a direct mapping between real tissue images and phenotype representations. We have illustrated the applicability of its representations in three different scenarios, providing evidence of the morphological and cellular information enfold on the real tissue representations. Furthermore, we found that these latent representations learn such stable and relevant features of histopathology images that they can be used as interpretable visualizations, as feature space of a linear classifier, or as input in a weakly supervised setting such as MIL.

We envision our model to be used as a tool to characterize phenotype patterns. These learned features of histopathology images can be used to find associations between tissue and matched clinical and biological data such as genomic, transcriptomic, or survival information, contributing to a further characterize tumor micro-environment and improved patient stratification.

Acknowledgements We will like to acknowledge funding support from University of Glasgow on A.C.Q scholarship, K.Y from EPSRC grant EP/R018634/1, and R.M-S. from EPSRC grants EP/T00097X/1 and EP/R018634/1. This work has used computing resources at the NYU School of Medicine High Performance Computing Facility.

References

1. Beck, A.H., Sangoi, A.R., Leung, S., Marinelli, R.J., Nielsen, T.O., van de Vijver, M.J., West, R.B., van de Rijn, M., Koller, D.: Systematic analysis of breast cancer morphology uncovers stromal features associated with survival. *Sci Transl Med* **3**(108), 108ra113 (Nov 2011). <https://doi.org/10.1126/scitranslmed.3002564>
2. Coudray, N., Ocampo, P.S., Sakellaropoulos, T., Narula, N., Snuderl, M., Fenyö, D., Moreira, A.L., Razavian, N., Tsirigos, A.: Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nature Medicine* **24**(10), 1559–1567 (2018). <https://doi.org/10.1038/s41591-018-0177-5>

3. Coudray, N., Tsirigos, A.: Deep learning links histology, molecular signatures and prognosis in cancer. *Nature Cancer* **1**(8), 755–757 (2020). <https://doi.org/10.1038/s43018-020-0099-2>
4. Fu, Y., Jung, A.W., Torne, R.V., Gonzalez, S., Vöhringer, H., Shmatko, A., Yates, L.R., Jimenez-Linan, M., Moore, L., Gerstung, M.: Pan-cancer computational histopathology reveals mutations, tumor composition and prognosis. *Nature Cancer* **1**(8), 800–810 (2020). <https://doi.org/10.1038/s43018-020-0085-8>
5. Hu, B., Tang, Y., Chang, E.I.C., Fan, Y., Lai, M., Xu, Y.: Unsupervised learning for cell-level visual representation in histopathology images with generative adversarial networks. *IEEE Journal of Biomedical and Health Informatics* **23**(3), 1316–1328 (May 2019). <https://doi.org/10.1109/jbhi.2018.2852639>
6. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: Dy, J., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 80, pp. 2127–2136 (2018)
7. Jolicoeur-Martineau, A.: The relativistic discriminator: a key element missing from standard GAN. In: *International Conference on Learning Representations* (2019)
8. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Jun 2019). <https://doi.org/10.1109/cvpr.2019.00453>
9. Kather, J.N., Halama, N., Marx, A.: 100,000 histological images of human colorectal cancer and healthy tissue (Apr 2018). <https://doi.org/10.5281/zenodo.1214456>
10. Kather, J.N., Heij, L.R., Grabsch, H.I., et al.: Pan-cancer image-based detection of clinically actionable genetic alterations. *Nature Cancer* **1**(8), 789–799 (2020). <https://doi.org/10.1038/s43018-020-0087-6>
11. Kather, J.N., Weis, C.A., Bianconi, F., Melchers, S.M., Schad, L.R., Gaiser, T., Marx, A., Zöllner, F.G.: Multi-class texture analysis in colorectal cancer histology. *Scientific Reports* **6**(1), 27988 (2016)
12. Krause, J., Grabsch, H.I., Kloor, M., Jendrusch, M., Echle, A., Buelow, R.D., Boor, P., Luedde, T., Brinker, T.J., Trautwein, C., Pearson, A.T., Quirke, P., Jenniskens, J., Offermans, K., van den Brandt, P.A., Kather, J.N.: Deep learning detects genetic alterations in cancer histology generated by adversarial networks. *The Journal of Pathology* **n/a**(n/a) (2021)
13. Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Xiaojun Guan, Schmitt, C., Thomas, N.E.: A method for normalizing histology slides for quantitative analysis. In: *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. pp. 1107–1110 (2009). <https://doi.org/10.1109/ISBI.2009.5193250>
14. Mahmood, F., Borders, D., Chen, R., McKay, G.N., Salimian, K.J., Baras, A., Durr, N.J.: Deep adversarial training for multi-organ nuclei segmentation in histopathology images. *IEEE Transactions on Medical Imaging* p. 1–1 (2020). <https://doi.org/10.1109/tmi.2019.2927182>
15. McInnes, L., Healy, J., Saul, N., Großberger, L.: UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software* **3**(29) (2018)
16. Quiros, A.C., Murray-Smith, R., Yuan, K.: PathologyGAN: Learning deep representations of cancer tissue. In: *Medical Imaging with Deep Learning* (2020)
17. Raczkowski, L., Mozejko, M., Zambonelli, J., Szczurek, E.: Ara: accurate, reliable and active histopathological image classification framework with bayesian deep learning. *Scientific Reports* **9**(1), 14347 (2019)

18. Rana, A., Yauney, G., Lowe, A., Shah, P.: Computational histological staining and destaining of prostate core biopsy RGB images with Generative Adversarial Neural Networks. 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA) (Dec 2018). <https://doi.org/10.1109/icmla.2018.00133>
19. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Computer Graphics and Applications* **21**(5), 34–41 (2001). <https://doi.org/10.1109/38.946629>
20. Schmauch, B., Romagnoni, A., Pronier, E., Saillard, C., Maillé, P., Calderaro, J., Kamoun, A., Sefta, M., Toldo, S., Zaslavskiy, M., Clozel, T., Moarii, M., Courtiol, P., Wainrib, G.: A deep learning model to predict rna-seq expression of tumours from whole slide images. *Nature Communications* **11**(1), 3877 (2020). <https://doi.org/10.1038/s41467-020-17678-4>
21. Woerl, A.C., Eckstein, M., Geiger, J., Wagner, D.C., Daher, T., Stenzel, P., Fernandez, A., Hartmann, A., Wand, M., Roth, W., Foersch, S.: Deep learning predicts molecular subtype of muscle-invasive bladder cancer from conventional histopathological slides. *European Urology* **78**(2), 256 – 264 (2020)
22. Yuan, Y., Failmezger, H., Rueda, O.M., Ali, H.R., Gräf, S., Chin, S.F., Schwarz, R.F., Curtis, C., Dunning, M.J., Bardwell, H., Johnson, N., Doyle, S., Turashvili, G., Provenzano, E., Aparicio, S., Caldas, C., Markowitz, F.: Quantitative image analysis of cellular heterogeneity in breast tumors complements genomic profiling. *Sci Transl Med* **4**(157), 157ra143 (Oct 2012). <https://doi.org/10.1126/scitranslmed.3004330>
23. Zanjani, F.G., Zinger, S., Bejnordi, B.E., van der Laak, J.A.W.M., de With, P.H.N.: Stain normalization of histopathology images using generative adversarial networks. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). pp. 573–577 (2018)

A Datasets

We provide here further details on the datasets and their processing:

1. The breast H&E cancer dataset is composed by the Netherlands Cancer Institute (NKI) and Vancouver General Hospital (VGH) cohorts with 248 and 328 patients, each of them with associated Tissue Micro-Arrays (TMAs). The original TMAs have a resolution of 1128×720 pixels, and we split each TMA into tiles of 224×224 , allowing them to overlap by 50%. We also perform data augmentation on these images, a rotation of 90° and 180° , and vertical and horizontal inversion. We filter out images in which the tissue covers less than 70% of the area. In addition, we label each tile based on the density of cancer cells in the image by using CRImage [22], a SVM classifier that quantifies the number of cancer cells, the number of other types of cells (such as stromal or lymphocytes), and the ratio of tumorous cells per area. We use nine different classes with class 8 accounting for tiles with the largest count of cancer cells. This dataset is composed by a training set of 249K tiles and 460 patients, and a test set of 13K tiles and 116 patients, with no overlapping patients between sets.
2. The colorectal H&E cancer dataset from National Center for Tumor diseases (NCT, Germany) [9] provides tissue images of 224×224 resolution with an associated type of tissue label: adipose, background, debris, lymphocytes, mucus, smooth muscle, normal colon mucosa, cancer-associated stroma, and colorectal adenocarcinoma epithelium (tumor). All tissue images are provided with stain normalization already applied [13]. The dataset is divided into a training set of 100K tissue patches and 86 patients, and a test set of 7K tissue patches and 50 patients, there is no overlapping patients between train and test sets. We use this dataset to visualize latent representations and reconstructions of the real tissue images, and to test the performance of a linear classifier over those same latent representations. Finally, in order to compare our model to supervised methods, we combine the classes stroma and smooth muscle into a class 'simple stroma' as in reference publications [11, 17].
3. The lung H&E cancer dataset contains samples with adenocarcinoma (LUAD), squamous cell carcinoma (LUSC), and normal tissue, composed by 1807 Whole Slide Images (WSIs) of 1184 patients from the Cancer Genome Atlas (TCGA). We make use of the pipeline provided in [2], diving each WSI into patches of 224×224 and filtering out images with less than 50% tissue in total area and apply stain normalization [19]. In addition, we label each slide as tumor and non-tumor depending on the presence of lung cancer in the tissue. Finally, we split the dataset into a training set of 916K tissue patches and 666 patients, and a test set of 569K tissue patches and 518 patients, with no overlapping patients between both sets. We use this dataset to apply multiple instance learning (MIL) over latent representations, testing the performance to predict the presence of tumor in the WSI.

B Visualization of tissue representations

In this section we provide additional examples of visualizations of breast and colorectal cancer tissue representations.

Fig. 5. Uniform Manifold Approximation and Projection (UMAP) vectors of tissue representations. We present breast cancer tissue from NKI and VGH, where images are labels based on cancer cell counts, class 8 accounting for the largest.

Fig.6. Uniform Manifold Approximation and Projection (UMAP) vectors of tissue representations. We present colorectal cancer tissue from NCT, where images are labeled based on the tissue type. We observe that real tissue with different morphology gets assigned distinct regions of the latent space.

C Tissue Reconstructions

Fig. 7. Real tissue images and their reconstructions. We take real tissue images and map them to the latent space with our encoder, then we use the generator with the latent vector representations to generate the image reconstructions. (a) correspond to the real tissue images X_{real} and (b) to the reconstructions $X_{recon} = G(E(X_{real}))$, the images are paired in columns. We present samples of colorectal cancer tissue from NCT with different tissue types: (1) colorectal adenocarcinoma epithelium (tumor), (2) cancer-associated stroma, (3) normal colon mucosa, (4) smooth muscle, (5) mucus, (6) lymphocytes, (7) debris, (8) background, and (9) adipose.

Fig. 8. Real tissue images and their reconstructions. We take real tissue images and map them to the latent space with our encoder, then we use the generator with the latent vector representations to generate the image reconstructions. (a) correspond to the real tissue images X_{real} and (b) to the reconstructions $X_{recon} = G(E(X_{real}))$, the images are paired in columns. We present samples of breast cancer tissue from VG and NKI.

D Logistic Regression Confusion Matrix and ROC

Fig.9. Confusion matrix and Receiver Operating Characteristic (ROC) curve of a logistic regression classifier for tissue type, trained over tissue representations.

E Network Architectures

Generator Network $G : w \rightarrow x$	
Dense Layer, adaptive instance normalization (AdaIN), and leakyReLU	$200 \rightarrow 1024$
Dense Layer, AdaIN, and leakyReLU	$1024 \rightarrow 12544$
Reshape	$7 \times 7 \times 256$
ResNet Conv2D Layer, 3x3, stride 1, pad same, AdaIN, and leakyReLU 0.2	$7 \times 7 \times 256 \rightarrow 7 \times 7 \times 256$
ConvTranspose2D Layer, 2x2, stride 2, pad upscale, AdaIN, and leakyReLU 0.2	$7 \times 7 \times 256 \rightarrow 14 \times 14 \times 512$
ResNet Conv2D Layer, 3x3, stride 1, pad same, AdaIN, and leakyReLU 0.2	$14 \times 14 \times 512 \rightarrow 14 \times 14 \times 512$
ConvTranspose2D Layer, 2x2, stride 2, pad upscale, AdaIN, and leakyReLU 0.2	$14 \times 14 \times 512 \rightarrow 28 \times 28 \times 256$
ResNet Conv2D Layer, 3x3, stride 1, pad same, AdaIN, and leakyReLU 0.2	$28 \times 28 \times 256 \rightarrow 28 \times 28 \times 256$
Attention Layer	at $28 \times 28 \times 256$
ConvTranspose2D Layer, 2x2, stride 2, pad upscale, AdaIN, and leakyReLU 0.2	$28 \times 28 \times 256 \rightarrow 56 \times 56 \times 128$
ResNet Conv2D Layer, 3x3, stride 1, pad same, AdaIN, and leakyReLU 0.2	$56 \times 56 \times 128 \rightarrow 56 \times 56 \times 128$
ConvTranspose2D Layer, 2x2, stride 2, pad upscale, AdaIN, and leakyReLU 0.2	$56 \times 56 \times 128 \rightarrow 112 \times 112 \times 64$
ResNet Conv2D Layer, 3x3, stride 1, pad same, AdaIN, and leakyReLU 0.2	$112 \times 112 \times 64 \rightarrow 112 \times 112 \times 64$
ConvTranspose2D Layer, 2x2, stride 2, pad upscale, AdaIN, and leakyReLU 0.2	$112 \times 112 \times 64 \rightarrow 224 \times 224 \times 32$
Conv2D Layer, 3x3, stride 1, pad same, 32 $\rightarrow$ 3	$224 \times 224 \times 32 \rightarrow 224 \times 224 \times 3$
Sigmoid	

Table 3. Generator Network Architecture details of Pathology GAN model.

Discriminator Network $C : x \rightarrow d$	
$x \in \mathbb{R}^{224 \times 224 \times 3}$	
ResNet Conv2D Layer, 3x3, stride 1, pad same, and leakyReLU 0.2	$224 \times 224 \times 3 \rightarrow 224 \times 224 \times 3$
Conv2D Layer, 2x2, stride 2, pad downscale, and leakyReLU 0.2	$224 \times 224 \times 3 \rightarrow 122 \times 122 \times 32$
ResNet Conv2D Layer, 3x3, stride 1, pad same, and leakyReLU 0.2	$122 \times 122 \times 32 \rightarrow 122 \times 122 \times 32$
Conv2D Layer, 2x2, stride 2, pad downscale, and leakyReLU 0.2	$122 \times 122 \times 32 \rightarrow 56 \times 56 \times 64$
ResNet Conv2D Layer, 3x3, stride 1, pad same, and leakyReLU 0.2	$56 \times 56 \times 64 \rightarrow 56 \times 56 \times 64$
Conv2D Layer, 2x2, stride 2, pad downscale, and leakyReLU 0.2	$56 \times 56 \times 64 \rightarrow 28 \times 28 \times 128$
ResNet Conv2D Layer, 3x3, stride 1, pad same, and leakyReLU 0.2	$28 \times 28 \times 128 \rightarrow 28 \times 28 \times 128$
Attention Layer at $28 \times 28 \times 128$	
Conv2D Layer, 2x2, stride 2, pad downscale, and leakyReLU 0.2	$28 \times 28 \times 128 \rightarrow 14 \times 14 \times 256$
ResNet Conv2D Layer, 3x3, stride 1, pad same, and leakyReLU 0.2	$14 \times 14 \times 256 \rightarrow 14 \times 14 \times 256$
Conv2D Layer, 2x2, stride 2, pad downscale, and leakyReLU 0.2	$14 \times 14 \times 256 \rightarrow 7 \times 7 \times 512$
Flatten $7 \times 7 \times 512 \rightarrow 25088$	
Dense Layer and leakyReLU, $25088 \rightarrow 1024$	
Dense Layer and leakyReLU, $1024 \rightarrow 1$	

Table 4. Discriminator Network Architecture details of Pathology GAN model.

Encoder Network $E : x \rightarrow w'$	
$x \in \mathbb{R}^{224 \times 224 \times 3}$	
Conv2D Layer, 2x2, stride 2, pad downscale, Instance Norm, and leakyReLU 0.2	$224 \times 224 \times 3 \rightarrow 224 \times 224 \times 32$
ResNet Conv2D Layer, 3x3, stride 1, pad same, Instance Norm, and leakyReLU 0.2	$224 \times 224 \times 32 \rightarrow 224 \times 224 \times 32$
Conv2D Layer, 2x2, stride 2, pad downscale, Instance Norm, and leakyReLU 0.2	$224 \times 224 \times 32 \rightarrow 122 \times 122 \times 64$
ResNet Conv2D Layer, 3x3, stride 1, pad same, Instance Norm, and leakyReLU 0.2	$122 \times 122 \times 64 \rightarrow 122 \times 122 \times 64$
Conv2D Layer, 2x2, stride 2, pad downscale, Instance Norm, and leakyReLU 0.2	$122 \times 122 \times 64 \rightarrow 56 \times 56 \times 128$
ResNet Conv2D Layer, 3x3, stride 1, pad same, Instance Norm, and leakyReLU 0.2	$56 \times 56 \times 128 \rightarrow 56 \times 56 \times 128$
Conv2D Layer, 2x2, stride 2, pad downscale, Instance Norm, and leakyReLU 0.2	$56 \times 56 \times 128 \rightarrow 28 \times 28 \times 256$
ResNet Conv2D Layer, 3x3, stride 1, pad same, Instance Norm, and leakyReLU 0.2	$28 \times 28 \times 256 \rightarrow 28 \times 28 \times 256$
Attention Layer at $28 \times 28 \times 256$	
Conv2D Layer, 2x2, stride 2, pad downscale, Instance Norm, and leakyReLU 0.2	$28 \times 28 \times 256 \rightarrow 14 \times 14 \times 512$
ResNet Conv2D Layer, 3x3, stride 1, pad same, Instance Norm, and leakyReLU 0.2	$14 \times 14 \times 512 \rightarrow 14 \times 14 \times 512$
Conv2D Layer, 2x2, stride 2, pad downscale, Instance Norm, and leakyReLU 0.2	$14 \times 14 \times 512 \rightarrow 7 \times 7 \times 512$
Flatten $7 \times 7 \times 512 \rightarrow 25088$	
Dense Layer and leakyReLU, $25088 \rightarrow 1024$	
Dense Layer and leakyReLU, $1024 \rightarrow 200$	

Table 5. Encoder Network Architecture details of Pathology GAN model.

Mapping Network $M : z \rightarrow w$
$z \in \mathbb{R}^{200} \sim \mathcal{N}(0, I)$
ResNet Dense Layer and ReLU, $200 \rightarrow 200$
ResNet Dense Layer and ReLU, $200 \rightarrow 200$
ResNet Dense Layer and ReLU, $200 \rightarrow 200$
ResNet Dense Layer and ReLU, $200 \rightarrow 200$
Dense Layer, $200 \rightarrow 200$

Table 6. Mapping Network Architecture details of Pathology GAN model.