

Attention-based Interpretable Regression of Gene Expression in Histology

Mara Graziani^{1,2}, Niccolò Marini², Nicolas Deutschmann¹, Nikita Janakaraian^{1,3}, Henning Müller², and María Rodríguez Martínez¹

¹ IBM Research Europe, 8803 Rüschlikon, Switzerland

² University of Applied Sciences of Western Switzerland, 3960 Sierre, Switzerland

³ ETH Zürich, 8092 Zürich, Switzerland

Abstract. Interpretability of deep learning is widely used to evaluate the reliability of medical imaging models and reduce the risks of inaccurate patient recommendations. For models exceeding human performance, e.g. predicting RNA structure from microscopy images, interpretable modelling can be further used to uncover highly non-trivial patterns which are otherwise imperceptible to the human eye. We show that interpretability can reveal connections between the microscopic appearance of cancer tissue and its gene expression profiling. While exhaustive profiling of all genes from the histology images is still challenging, we estimate the expression values of a well-known subset of genes that is indicative of cancer molecular subtype, survival, and treatment response in colorectal cancer. Our approach successfully identifies meaningful information from the image slides, highlighting hotspots of high gene expression. Our method can help characterise how gene expression shapes tissue morphology and this may be beneficial for patient stratification in the pathology unit. The code is available on GitHub.

Keywords: interpretability · histopathology · transcriptomics · attention

1 Introduction

The wide variability of existing interpretability methods, in the form of post-hoc explanations [1,2] or models with ad-hoc transparency constraints [3], has been mainly dedicated to ensuring the safety and reliability of opaque deep neural networks. In medical applications such as digital pathology, saliency maps highlighted the relevance of anomalous nuclei in the detection of tumorous tissue [4], and concept-based analyses confirmed that clinically relevant measures on nuclei area and appearance are learned as intermediate features [5]. Undesired hidden biases and behaviours were detected and corrected to improve model performance [6,7,8]. Only recently, interpretability techniques were proposed to uncover unknown insights about models with super-human performance, e.g. the algorithm AlphaZero [9] defeating the world chess master. Similarly, deep learning models can predict biomarkers invisible to the human eye [10,11]. For

instance, DNA mutations [10] and gene expression profiles were inferred from hematoxylin and eosin (H&E) stained whole slide images (WSIs) [11], demonstrating that complex transcriptomic patterns can be captured from images without the support of sophisticated sequencing machines.

While ensuring patient safety remains the top priority in high-risk clinical environments [12,13], scientific discovery may benefit from unravelling the complex association between RNA expression and tissue microscopy. Identifying the histological patterns that are predictive of gene expression could have an important impact on diagnostic routines, facilitating early patient stratification and aiding the identification of molecular subtypes that are informative of prognosis, survival, and treatment response [14]. For example, clinical trials show that some colorectal (CRC) cancer patients with hypermutated microsatellite instability (MSI) may respond better to immunotherapy than chemotherapy [15].

Our central question is whether we can clarify which regions in CRC H&E slides are the most informative about RNA transcriptomics by using a trainable attention mechanism [16]. It is yet unclear whether the existing algorithms [11,17] infer the bulk average expression of RNA from hotspots of high expression or as a uniform distribution on the entire slide. To obtain interpretable insights, we learn the expression of individual genes rather than the entire transcriptomic profile at once, training several attention-based multiple instance regression models independently. This enables a fine grained analysis of each gene individually and reduces the need for extremely large dataset sizes [11,17]. In our results, the attention mechanism brings transparency to the morphological patterns in the tissue that are learned by the model. Hotspots of high expression are highlighted in most of our visualisations and hypotheses can be formulated to relate the histological appearance of tissue to varying gene expression. We show that meaningful information is successfully filtered from the WSIs by the attention, leading to more accurate gene expression estimates and patient stratification. A reduction of 10% in the regression error is seen across genes and patients.

2 Methods

2.1 Datasets

We use images of colon adenocarcinoma (COAD) and rectal adenocarcinoma (READ) that are publicly available together with matched transcriptomic profiles at The Cancer Genome Atlas (TCGA)⁴. Each biopsy sample is split into three adjacent portions, two of which are used to generate H&E-stained frozen tissue slides, called top-section (TS) and bottom-section (BS), used to verify the presence of sufficient tumour content before sequencing. The central section is used for RNA sequencing. Differently from [11,17], we focus on frozen tissue sections rather than on diagnostic slides, since they are directly adjacent to the sequenced tissue, and thus constitute the best available representation of the

⁴ <https://portal.gdc.cancer.gov>, as accessed in June 2022.

RNA profiles [18]. The WSIs are preprocessed with HistoQC [19] to mask out the background and blurred locations.

Gene expression profiles are obtained from the UCSC Xena Browser [20], which links to the Genomic Data Commons [21] version of the TCGA COAD and READ projects. The High Throughput-Sequencing (HT-Seq) raw counts are log2-based Fragments per Kilobase of transcript per Million mapped reads (FPKM) normalised. Because of a distributional shift among institutions, we retain all the patient measurements from a single institution, ignoring duplicates for 23 patients. We focus on the 45 genes in the ColoType signature [22], a gene set that is predictive of CRC prognosis [23] and clinico-pathological variables [24]. The list of genes is in the Appendix Table A.1. Single-cell RNA sequencing data are also used for validation⁵. Namely, we use the profiling of 969 single-cells from the CRC resected primary tumours of 11 patients in [25].

As in [26], we exclude patients in preoperative therapy and the rare subtypes of neuroendocrine and signet cell tumours. In total, we use 774 WSIs at 20X magnification from 364 patients. The test set is built by selecting randomly 82 patients and the remaining 282 patients are used for five-fold cross validation.

2.2 Multiple Instance Regression of Gene Expression




Fig. 1: Attention-based Multiple Instance Learning for the semi-transparent regression of RNA gene expression from H&E WSIs.

Our model is an adaptation of the attention-based multiple instance learning approach in [16], which follows the implementation in [27]. We consider a patient p with associated H&E-stained scans of the top and bottom frozen tissue sections x_t and x_b , both images in $\mathbb{R}^{w \times h}$, where w is the image width and h is the height. The associated bulk-RNA expression for a given gene g is $y_g \in \mathbb{R}$. To deal with the large input sizes of WSIs, x_t and x_b are represented as a single collection of N non-overlapping adjacent patches obtained by a sliding window $\{x_i\}_{i=0}^N$ [28]. The task is to learn a permutation invariant mapping to the real-valued gene label y_g from this joint collection. We assume that each collection contains at least one instance associated with the label [29]. The gene expression y_g can be obtained as a linear-weighted combination of the $y_{i,g}$ predictions for each of the x_i patches,

⁵ The dataset is available at <https://www.weizmann.ac.il/sites/3CA/colorectal>

where the weights are given by the trainable attention. Differently from [16,27], the attention weights $\{a_{i,g}\}_{i=0}^N$ are optimised to predict the continuous gene expression label rather than a binary outcome.

The model, illustrated in Figure 1, comprises: (i) a convolutional backbone to obtain low-dimensional representations of the input patches, i.e. a ResNet18 [30] pretrained on ImageNet [31]; (ii) a dense layer to predict the gene expression $y_{i,g}$ for each input x_i ; (iii) an attention network to learn the attention weights.

The output of our model is the expression of a single gene, and we train a different model for each gene in our selection. This is opposed to the approach of the SoA [11], where regression values for over 30000 gene outputs are optimised in a single training. To allow for a fair comparison with the SoA, we adapt their model to output the expression of a single gene at a time. The adapted model predicts the gene expression for each patch, and the predictions are then aggregated by a weighted average with larger weights given to large-valued patch predictions as in [11]. Our trainable attention removes the need for this heuristic.

2.3 Attention-based Model Interpretability

The trainable attention provides transparency on how the model filters the information in the WSIs. Salient regions are identified without the need for an explicit localisation module [16,32]. However, recent debate argued that multiple plausible attention-based explanations may exist and that attention should be interpreted carefully [33,34]. The research in [35] addressed that debate, mentioning that attention weights learned by multiple models should be ensembled by either max or average pooling to reduce the risks of obtaining misleading interpretations. We thus average the attention weights of the models trained on the five folds, ensuring a trustworthy interpretation that focuses only on signal-bearing instances.

2.4 Evaluation of performance and interpretability

Model performance is evaluated in multiple ways. First, the Median Average Percentage Error (MAPE) of our predictions is compared against the SoA. Both models are also compared to the lower bound of random guessing given by predicting the mean of the training labels. Finally, the significance of our results is verified by computing the non-parametric one-tailed Wilcoxon test between the patient-wise percentage error made by our model and the SoA. Successfully regressed genes are identified as in [11] by evaluating the Pearson’s correlation coefficient ρ between the true and the predicted labels.

The evaluation of model interpretability is non-trivial [17,35]. For morphological patterns that are known to be associated with RNA abnormal expression, e.g. mucinous tissue showing high levels of *MUC2*, we verify that our explanations confirm the existing knowledge. For the newly identified patterns, however, there is little to no ground truth available. In this case, we evaluate how well patch-wise attention weights correlate to the ground truth obtained from analysing single-cell data. We identify, for instance, genes that are co-expressed




Fig. 2: Median difference between the test set errors by our method and the SoA. Statistical significance is shown for p-value < 0.05. Lower is better.

in single-cell data according to their Pearson’s correlation, and we verify that the patch-wise attentions in overlapping WSI locations correlate for those models. Finally, the benefits of the attention are quantified by quantifying the performance improvement on an auxiliary task, i.e. inferring MSI status. The higher the quality of the attention-based localisation, the more accurate we expect to be the MSI-based stratification of the patients.

3 Experiments and Results

3.1 Network Training

The models are trained by optimising the mean squared error loss between the predicted gene expression value \hat{y}_g and the label y_g . Stochastic gradient descent (SGD) is used with standard hyperparameters (learning rate 0.0001, momentum 0.9, weight decay 0.01) and early stopping (12 epochs patience). The gradient updates affect only the dense layers and the attention mechanism. MSI status is learned in Sec 3.4 by two additional dense layers trained by SGD minimisation of the weighted binary cross-entropy. For a single model trained on a GPU Tesla V100 for 5 hours, we estimate a carbon footprint of 0.65 kgCO₂ [36].

3.2 Quantitative Model Evaluation

The gene expression estimates obtained by our method are more accurate than the SoA. Across all genes and patients, our model obtains a reduction of 10% in the MAPE, i.e. 0.65 ± 0.08 in our model against 0.72 ± 0.12 in the SoA. Random guessing as described in Sec. 2.4 achieves MAPE 0.94 ± 13.0 , showing that both methods are capturing meaningful signals from the WSIs. Computing the Pearson’s correlation with the ground-truth labels further confirms the accuracy of our model. *MGP* achieves the highest ρ with 0.71 against 0.69 of the SoA (p-value < 0.0001). The predictions for *CCDC80*, *NRP2*, and *RAB34* also show strong correlation with the ground-truth, with ρ 0.65, 0.62 and 0.61 against the SoA ρ at 0.61, 0.61 and 0.55 respectively. The median differences in the errors made by the two models are shown in Figure 2. The statistical significance of the Wilcoxon test is reported for the individual patient-wise error differences. The detailed results on the full gene set are provided in the Appendix Figure B.1 and Table B.1.




Fig. 3: Localised ensemble attention on an unseen test image for six genes.

The largest improvement is observed for *SPINK4*, where the MAPE decreases to 0.62 ± 1.9 from 0.88 ± 3.1 of the SoA (p-value < 0.001). *QPRT*, *MUC2*, *AOC3* and *SDC2* report significantly better MAPE than the SoA, achieving respectively 0.74 ± 1.3 , 0.73 ± 1.2 , 0.55 ± 1.3 and 0.64 ± 1.4 against 0.98 ± 1.1 , 0.95 ± 1.9 , 0.65 ± 3.7 and 0.81 ± 2.3 (p-value < 0.05). The Pearson’s ρ for *QPRT* and *MUC2*, in particular, increase respectively to 0.43 and 0.46 (p-value < 0.0001) from 0.35 and 0.13 for the SoA.

3.3 Attention-based Identification of Hotspots and Patterns

Figures 3 and 4 show the predictions and the distribution of the attention for our best gene models on the input WSI. The attention identifies hotspot regions, rather than being uniformly distributed over all patches. Figure 5 visualises more in detail the morphological patterns in the highlighted hotspots of high and low gene expression. We retrieve, for instance, the patches that received the highest normalised attention weights, namely the highest $N_p a_{i,g} \forall i, p$ in the testing set, where $i = 1, \dots, N_p$, and N_p is the total number of patches for patient p ⁶. Additional visualisations are in the Appendix and the code repository⁷.

Genes that are co-expressed in single-cell data show high attention on the same WSI areas. For example, *NRP2* and *MUC2* show a correlation value of $\rho = 0.40$ at the single-cell gene expression level. Similarly, they display a correlation of 0.63 in the patch-wise attention values. Conversely, *MUC2* is not co-expressed at the single-cell level with *AOC3* and *COL8A2*, and the attention localise on

⁶ The normalisation enables the comparison across patients, since $\sum_{i=0}^{N_p} a_{i,g} = 1$




Fig. 4: Left-to-right order: Original input WSI, raw spatial predictions and attention-weighted predictions. Blue and red show low and high expression respectively.




Fig. 5: Patches with highest attention for predicted low (blue box) and high (red box) expression. Differing patterns are visible across genes and expression levels.

different regions, with 0.16 and 0.12 patch-wise correlations respectively. The results for all gene couples are in the Appendix Figure [D.1](#)

3.4 Quantitative Evaluation of the Attention

The MSI status is inferred from the gene expression values regressed from the H&E slides. The input noise is reduced by selecting only the genes that were predicted with Pearson’s p -value < 0.05 by each method. This leads to 37 genes with our method and 35 with the SoA from the initial pool of 45 genes. We perform 5x2 cross-validation to obtain more stable estimates through repetitions on non-overlapping training sets [37](#). We observe an increase in the AUC from 0.60 ± 0.08 of the SoA to 0.71 ± 0.05 with our attention-based model. In comparison, using the ground-truth sequencing data yields AUC 0.86 ± 0.04 .

4 Discussion

The experiments in Sec. [3.2](#) quantify the accuracy of our bulk-RNA predictions from WSIs in terms of the prediction error, i.e. MAPE, and correlation with the ground-truth labels, i.e. Pearson correlation ρ . Our trainable attention mechanism reduces the error for the majority of the analysed genes, achieving significant improvements over the SoA for the *SPINK4*, *QPRT*, *MUC2*, *SDC2* and *AOC3* genes. Besides accuracy improvements, learning the attention removes the use of a priori heuristic choices and hyperparameter tuning in [11](#).

The main contribution is the analysis of the signals that are picked by the models to obtain our performance improvements. Focusing on one gene at a time enables a fine-grained analysis of the model attention patterns since the entire model capacity is optimised to capture salient features of individual genes. The attention maps in Sec. [3.3](#) give interesting insights. For instance, Figure [3](#) shows that the attention localises over hotspot regions rather than being uniformly spread across the slide. This answers the questions raised in [17](#) regarding the distribution of the information related to gene expression. As expected, high attention is seen in overlapping WSI regions for genes that are co-expressed in the same cell and in differing regions for genes that are not co-expressed. Moreover, the highest attention weights point to visibly different histological

patterns, as shown in Figure 5. For example, the high-expression patches of *COL8A2* present higher concentrations of stroma than the low-expression ones.

An interesting example is that of *MUC2*, a particularly well-known gene that impacts tissue histology. *MUC2* over-expression leads to high amounts of visible intra- and extra-cellular mucin, characterizing the tissue as mucinous adenocarcinoma. An example of mucinous adenocarcinoma is shown in Figure 4b. Our model predicts high values of *MUC2* for the regions in the slides where the mucinous tissue is the most visible. When these are further weighted by the attention, the localisation is highly focused on a few hotspots of high expression. From a quantitative standpoint, the impact of the attention mechanism for this specific gene is remarkably significant. The SoA predictions for *MUC2* are similar to random guessing (with MAPE 0.95 and ρ 0.13), whereas our method significantly reduces the MAPE to 0.73 and the correlation with ground truth values, i.e. ρ , to 0.46 with p-value < 0.0001. This result is also beneficial to the prediction of MSI, which is known to correlate with the presence of mucin [26].

5 Conclusion

We proposed an attention-based multiple instance regression model to infer bulk gene expression of tissue sections from H&E histology slides of colorectal adenocarcinoma. The automated regression of transcriptomics from H&E is not yet a feasible replacement for sequencing the tissue, but the analysis of interpretability patterns can highlight some histological patterns associated with specific gene expression levels. We show here how the attention mechanism successfully filters informative content from the vast amount of information in the WSIs, leading to significantly lower errors and interesting insights into how gene expression impacts tissue morphology. Further developments of this method should introduce pathologists in the loop to validate the identified patterns and further investigate their association to differentially expressed genes. In turn, this could help pathologists to identify cancer molecular subtypes from WSIs and consequently stratify patients for targeted treatment.

Acknowledgements This work was supported by the Swiss National Science Foundation Sinergia project (CRSII5_193832) and the EU H2020 project AI4Media (951911).

References

1. B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *CVPR 2016*, pages 2921–2929, 2016.
2. Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *KDD 2016*, pages 1135–1144, 2016.
3. C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.

4. M. Graziani, I. Palatnik de Sousa, M. M. B. R. Vellasco, E. Costa da Silva, H. Müller, and V. Andrearczyk. Sharpening local interpretable model-agnostic explanations for histopathology: Improved understandability and reliability. In *MICCAI*, 2021.
5. M. Graziani, V. Andrearczyk, S. Marchand-Maillet, and H. Müller. Concept attribution: Explaining CNNs to physicians. *CBM*, 123:103865, 2020.
6. M. Hägele, P. Seegerer, S. Lapuschkin, M. Bockmayr, W. Samek, F. Klauschen, K. Müller, and A. Binder. Resolving challenges in deep learning-based analyses of histopathological images using explanation methods. *Scientific reports*, 10(1):6423, April 2020.
7. M. Graziani, T. Lompech, H. Müller, A. Depeursinge, and V. Andrearczyk. On the Scale Invariance in State of the Art CNNs Trained on ImageNet. *MAKE*, 3(2):374–391, 2021.
8. M. Graziani, S. Otálora, S. Marchand-Maillet, H. Müller, and V. Andrearczyk. Learning Interpretable Pathology Features by Multi-task and Adversarial Training Improves CNN Generalization. *arXiv:2008.01478*, 2021.
9. T. McGrath, A. Kapishnikov, N. Tomasev, A. Pearce, D. Hassabis, B. Kim, U. Paquet, and V. Kramnik. Acquisition of Chess Knowledge in AlphaZero. *ArXiv:2111.09259*, 2021.
10. M. Chen, B. Zhang, W. Topatana, J. Cao, H. Zhu, S. Juengpanich, Q. Mao, H. Yu, and X. Cai. Classification and mutation prediction based on histopathology H&E images in liver cancer using deep learning. *NPJ precision oncology*, 4(1):1–7, 2020.
11. B. Schmauch, A. Romagnoni, E. Pronier, C. Saillard, P. Maillé, J. Calderaro, A. Kamoun, M. Sefta, S. Toldo, M. Zaslavskiy, et al. A deep learning model to predict rna-seq expression of tumours from whole slide images. *Nature communications*, 11(1):1–15, 2020.
12. Bas H.M., Hugo J. K., Kenneth G.A. G., and Max A. V. Explainable artificial intelligence (XAI) in deep learning-based medical image analysis. *MIA*, 79, 2022.
13. B. J. Lengerich, R. Caruana, M. E. Nunnally, and M. Kellis. Death by round numbers and sharp thresholds: How to avoid dangerous ai ehr recommendations. *medRxiv*, 2022.
14. J. Guinney, R. Dienstmann, X. Wang, A. De Reynies, A. Schlicker, C. Sonesson, L. Marisa, P. Roepman, G. Nyamundanda, P. Angelino, et al. The consensus molecular subtypes of colorectal cancer. *Nature medicine*, 21(11):1350–1356, 2015.
15. A. Trullas, J. Delgado, A. Genazzani, J. Mueller-Berghaus, S. Migali, C. Müller-Egert, H. Zander, H. Enzmann, and F. Pignatti. The ema assessment of pembrolizumab as monotherapy for the first-line treatment of adult patients with metastatic microsatellite instability-high or mismatch repair deficient colorectal cancer. *ESMO Open*, 6(3):100145, 2021.
16. M. Ilse, J. Tomczak, and M. Welling. Attention-based deep multiple instance learning. In *ICML*, pages 2127–2136. PMLR, 2018.
17. P. Weitz, Y. Wang, J. Hartman, and M. Rantalainen. An investigation of attention mechanisms in histopathology whole-slide-image analysis for regression objectives. In *ICCV*, pages 611–619, 2021.
18. L. AD. Cooper, E. G. Demicco, J. H. Saltz, R. T. Powell, A. Rao, and A. J. Lazar. Pancancer insights from the cancer genome atlas: the pathologist’s perspective. *The Journal of pathology*, 244(5):512–524, 2018.
19. A. Janowczyk, R. Zuo, H. Gilmore, M. Feldman, and A. Madabhushi. Histoqc: an open-source quality control tool for digital pathology slides. *JCO clinical cancer informatics*, 3:1–7, 2019.

20. M. J. Goldman, B. Craft, M. Hastie, K. Repečka, F. McDade, A. Kamath, A. Banerjee, Y. Luo, D. Rogers, et al. Visualizing and interpreting cancer genomics data via the Xena platform. *Nature biotechnology*, 38(6):675–678, 2020.
21. R. L. Grossman, A. P. Heath, V. Ferretti, H. E. Varmus, D. R. Lowy, W. A. Kibbe, and L. M. Staudt. Toward a shared vision for cancer genomic data. *New England Journal of Medicine*, 375(12):1109–1112, 2016.
22. S. A. Buechler, M. T. Stephens, A. B. Hummon, K. Ludwig, E. Cannon, T. C. Carter, J. Resnick, Y. Gökmen-Polar, and S. S. Badve. Colotype: a forty gene signature for consensus molecular subtyping of colorectal cancer tumors using whole-genome assay or targeted rna-sequencing. *Scientific reports*, 10(1):1–13, 2020.
23. F. Pan, T. Chen, X. Sun, K. Li, X. Jiang, A. Försti, Y. Zhu, and M. Lai. Prognosis prediction of colorectal cancer using gene expression profiles. *Frontiers in Oncology*, 9, 2019.
24. E. AH. Kheirleiseid, N. Miller, K.H. Chang, M. Nugent, and M. J. Kerin. Clinical applications of gene expression in colorectal cancer. *Journal of gastrointestinal oncology*, 4(2):144, 2013.
25. H. Li, E. T. Courtois, D. Sengupta, Y. Tan, K.H. Chen, J.J.L. Goh, S.L. Kong, C. Chua, L.K. Hon, W.S. Tan, et al. Reference component analysis of single-cell transcriptomes elucidates cellular heterogeneity in human colorectal tumors. *Nature genetics*, 49(5):708–718, 2017.
26. H.G. Nguyen, O. Lundström, A. Blank, H. Dawson, A. Lugli, M. Anisimova, and I. Zlobec. Image-based assessment of extracellular mucin-to-tumor area predicts consensus molecular subtypes (cms) in colorectal cancer. *Modern pathology*, pages 1–9, 2021.
27. N. Marini, S. Marchesin, S. Otálora, M. Wodzinski, A. Caputo, M. van Rijthoven, W. Aswolinskiy, J. Bokhorst, D. Podareanu, E. Petters, et al. Unleashing the potential of digital pathology data by training computer-aided diagnosis models without human annotations. *npj Digital Medicine*, 5(1):1–18, 2022.
28. N. Marini, S. Otálora, D. Podareanu, M. van Rijthoven, J. van der Laak, F. Ciompi, H. Müller, and M. Atzori. Multi_Scale_Tools: a python library to exploit multi-scale WSIs. *Data-Enabled Intelligence for Medical Technology Innovation*, VI, 2022.
29. R. Soumya and P. David. Multiple Instance Regression. In *ICML*, pages 425–432. Morgan Kaufmann, 2001.
30. K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
31. O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
32. J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, and D. Rueckert. Attention gated networks: Learning to leverage salient regions in medical images. *MIA*, 53:197–207, 2019.
33. S. Jain and B. C. Wallace. Attention is not explanation. In *Proceedings of NAACL-HLT*, pages 3543–3556, 2019.
34. S. Wiegrefe and Y. Pinter. Attention is not not explanation. In *EMNLP-IJCNLP*, pages 11–20, Hong Kong, China, November 2019. Association for Computational Linguistics.
35. J. Haab, N. , Deutschmann, and M.R. Martínez. Is Attention Interpretation? A Quantitative Assessment on Sets. In *arXiv preprint arXiv:2207.13018*, 2022.
36. A. Lacoste, A. Luccioni, V. Schmidt, and T. Dandres. Quantifying the carbon emissions of machine learning. *arXiv:1910.09700*, 2019.

37. E. Alpaydmn. Combined 5×2 CV F Test for Comparing Supervised Classification Learning Algorithms. *Neural Computation*, 11 (8):1885–1892, 1999.

A Description of Selected Genes

Table A.1: Summary of the 45 genes considered in this study, including the ColoType signature and biomarkers of colorectal adenocarcinoma.

Gene	Expression level	Associated CMS	Ref.	Description
ASPHD2	HIGH	CMS1	[22]	Dioxygenase activity
ATP9A	LOW	CMS1	[22]	Membran trafficking of cargo proteins
AXIN2	LOW	CMS1	[22]	Regulation of beta-catenin stability in WNT pathway
CDHR1	LOW	CMS1	[22]	Cadherin superfamily, cell adhesion
CTTNBP2	LOW	CMS1	[22]	Cortacting binding protein
DACH1	LOW	CMS1	[22]	Chromatin associated protein
GNLY	HIGH	CMS1	[22]	Antimicrobial protein that kills intracellular pathogens.
HPSE	HIGH	CMS1	[22]	Enhances angiogenesis
SEMA5A	LOW	CMS1	[22]	Promotes angiogenesis by cell proliferation and migration and inhibites apoptosis
WARS	HIGH	CMS1	[22]	Regulates ERK, Akt, and eNOS pathways, associated with angiogenesis
CEL	HIGH	CMS2	[22]	
DDX27	HIGH	CMS2	[22]	Probable ATP-dependent RNA helicase
DUSP4	LOW	CMS2	[22]	Regulates mitogenic signal transduction
FSCN1	LOW	CMS2	[22]	Organizes filamentous actin into parallel bundles
LYZ	LOW	CMS2	[22]	
PLAGL2	HIGH	CMS2	[22]	
POFUT1	HIGH	CMS2	[22]	
QPRT	HIGH	CMS2	[22]	
TP53RK	HIGH	CMS2	[22]	
TRIB2	LOW	CMS2	[22]	
ASRGL1	HIGH	CMS3	[22]	
B3GNT6	HIGH	CMS3	[22]	Plays an important role in the synthesis of mucin-type O-glycans in digestive organs
CAPN9	HIGH	CMS3	[22]	
FBN1	LOW	CMS3	[22]	
FCGBP	HIGH	CMS3	[22]	
RASD1	HIGH	CMS3	[22]	
RBMS1	LOW	CMS3	[22]	
SPINK4	HIGH	CMS3	[22]	
TIMP3	LOW	CMS3	[22]	
VAV2	LOW	CMS3	[22]	
AOC3	HIGH	CMS4	[22]	Participates in lymphocyte extravasation and recirculation
ARMCX1	HIGH	CMS4	[22]	Regulates mitochondrial transport during axon regeneration
CCDC80	HIGH	CMS4	[22]	Promotes cell adhesion and matrix assembly
COL8A2	HIGH	CMS4	[22]	Necessary for migration and proliferation of vascular smooth muscle cells
MGP	HIGH	CMS4	[22]	Thought to act as inhibitor of bone formation
NRP2	HIGH	CMS4	[22]	May play a role in cardiovascular development, axon guidance, and tumorigenesis
RAB34	HIGH	CMS4	[22]	
SDC2	HIGH	CMS4	[22]	participates in cell proliferation, cell migration and cell-matrix interactions
TGFB3	HIGH	CMS4	[22]	Involved in embryogenesis and cell differentiation
TNS1	HIGH	CMS4	[22]	
DNMT3B	UNK	UNK	[23]	Required for genome-wide de novo methylation. Seems to be involved in gene silencing
CDKN2A	UNK	UNK	[24]	Many studies suggest poorer prognostic outcome for patients with hypermethylation in colorectal, liver, and younger lung cancer patients
CEACAM5	UNK	UNK	[24]	Cell adhesion, intracellular signaling and tumor progression
CXCR4	UNK	UNK	[24]	Essential for the vascularization of the gastrointestinal tract
MUC2	UNK	UNK	[24]	Coats the epithelia in the colon. May exclude bacteria from the inner mucus layer.

B Detailed Model Evaluation




Fig.B.1: Absolute Percentage Error on the test set for the baseline model (in blue) and our method (in green). The models are trained on five training folds and the final prediction is obtained as the ensemble average of the predictions. The red horizontal line shows the performance of random prediction. The lower the better.

Detailed results are reported in Table [B.1](#), which was not included in the main paper because of the limited space. The model of the gene expression for *NRP2* leads to the best results overall, with $\text{MAPE}=0.47 \pm 1.19$ and $\rho = 0.62$ against the SoA $\text{MAPE}=0.51 \pm 1.17$ and $\rho = 0.61$ of the baseline.

C Additional Visualizations

Additional visualizations of the spatialized prediction maps with and without attention-based weighting are given in Figure [C.1](#). Figures [C.2](#) and [C.3](#) show in detail the histologic patterns selected by the attention for the genes *COL8A2* and *MUC2* respectively.

D Single-cell co-expression

The co-expression of genes in single-cell RNA sequencing data is summarized in Figure [D.2](#). Figure [D.1](#) compares the correlation of the attention weights in overlapping patches to the co-expression of genes in single-cell data in terms of their Pearson’s correlation.

The couples (i) *CCDC80*, *TIMP3*; (ii) *CCDC80*, *MGP*; (iii) *TIMP3*, *MGP*; (iv) *FCGBP*, *MUC2* report high correlation with Pearson’s ρ being 0.85, 0.77, 0.69, 0.62 respectively. These couples show high correlation in the attention patterns. Figure [D.3](#) shows the similarity in the attention patterns for *FCGBP* and *MUC2*.

Table B.1: Comparison of the Test MAPE. Standard deviation reported in brackets.

Gene	MAPE ↓	
	SoA	ours
NRP2	0.51 (1.2)	0.47 (1.2)
COL8A2	0.54 (4.1)	0.49 (3.5)
CEACAM5	0.59 (0.8)	0.52 (0.9)
WARS	0.55 (6.5)	0.52 (12.3)
CCDC80	0.56 (2.0)	0.54 (1.5)
TNS1	0.54 (1.3)	0.54 (1.0)
ASPHD2	0.58 (1.9)	0.55 (2.7)
AOC3	0.65 (3.7)	0.55 (1.3)
POFUT1	0.70 (0.9)	0.55 (1.4)
TGFB3	0.67 (6.2)	0.57 (2.7)
CXCR4	0.65 (3.2)	0.60 (4.1)
TIMP3	0.66 (3.4)	0.60 (2.6)
VAV2	0.74 (1.3)	0.61 (1.3)
FBN1	0.58 (1.6)	0.61 (1.9)
PLAGL2	0.74 (1.0)	0.61 (1.4)
B3GNT6	0.76 (5.0)	0.62 (2.8)
SPINK4	0.88 (3.1)	0.62 (1.9)
AXIN2	0.71 (0.7)	0.63 (0.7)
DNMT3B	0.65 (9.3)	0.64 (10.5)
SDC2	0.81 (2.3)	0.64 (1.4)
MGP	0.64 (1.6)	0.65 (3.0)
SEMA5A	0.80 (1.7)	0.65 (1.4)
RBMS1	0.72 (1.1)	0.65 (1.7)
RASD1	0.69 (5.8)	0.66 (4.5)
APC	0.71 (2.1)	0.67 (2.0)
GNLY	0.62 (4.8)	0.67 (3.5)
ASRGL1	0.74 (4.1)	0.68 (4.1)
HPSE	0.71 (6.7)	0.68 (5.0)
CDKN2A	0.67 (36.1)	0.69 (29.8)
CAPN9	0.82 (5.1)	0.69 (2.6)
RAB34	0.69 (2.3)	0.69 (5.5)
ATP9A	0.64 (0.9)	0.70 (1.0)
FSCN1	0.62 (5.0)	0.70 (4.8)
TP53RK	0.74 (1.2)	0.71 (0.9)
CTTNBP2	0.76 (3.0)	0.71 (2.2)
DDX27	0.68 (1.0)	0.72 (2.1)
CEL	0.77 (1.2)	0.72 (1.4)
ARMCX1	0.66 (15.1)	0.73 (15.1)
MUC2	0.95 (1.9)	0.73 (1.2)
TP53	0.76 (0.9)	0.74 (0.9)
DUSP4	0.78 (3.1)	0.74 (3.5)
QPRT	0.98 (1.1)	0.74 (1.3)
CDHR1	0.87 (1.2)	0.76 (1.2)
FCGBP	0.89 (2.5)	0.77 (2.5)
DACH1	1.02 (2.9)	0.90 (2.2)

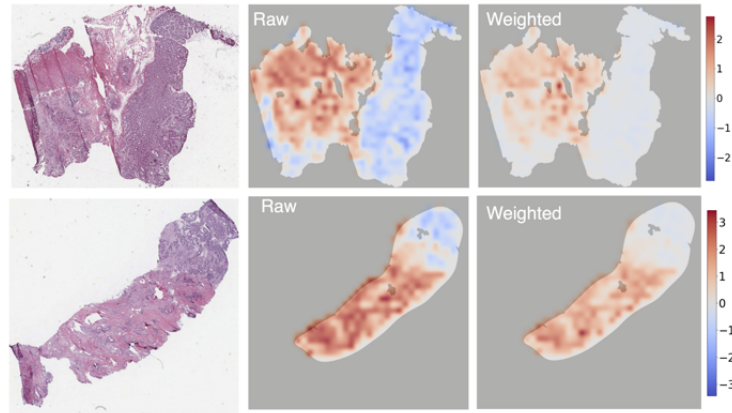


Fig. C.1: Additional spatialized prediction maps for *AOC3*




Fig. C.2: High-attention patches against gene expression levels (gex) for *COL8A2*

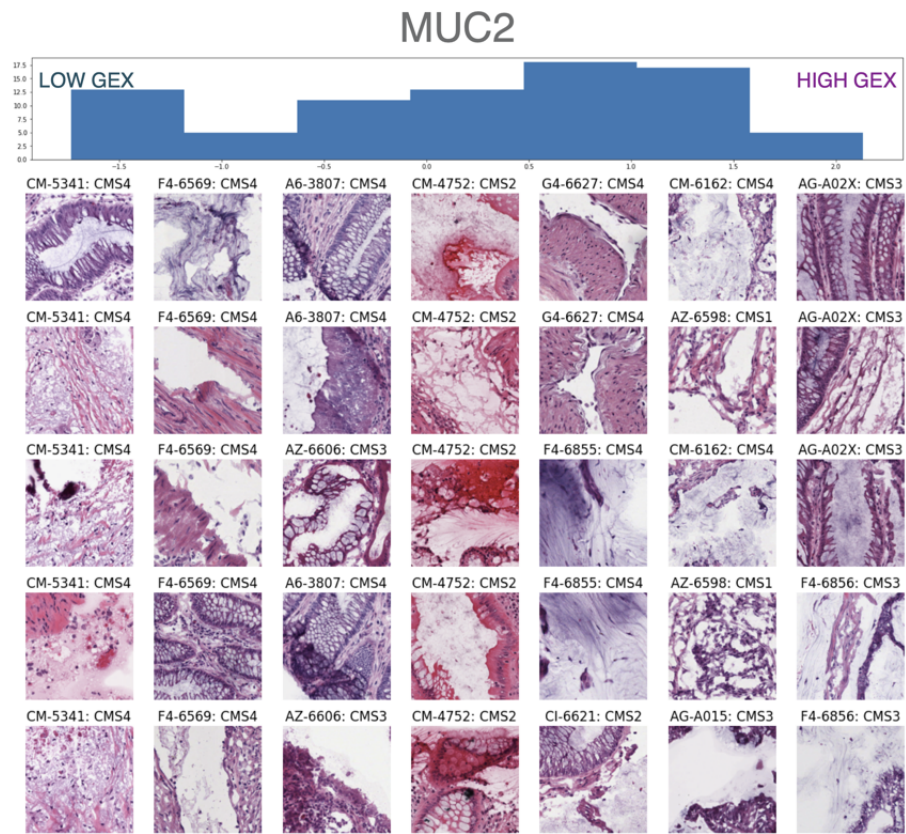


Fig. C.3: High-attention patches against gene expression levels (gex) for *MUC2*




Fig. D.1: Correlation of attention patterns in overlapping WSIs regions for co-expressed genes with correlation p-value < 0.05.




Fig. D.2: Summary of gene co-expression in single-cells. Each cell in the matrix represents the Pearson's correlation of the gene expressions.

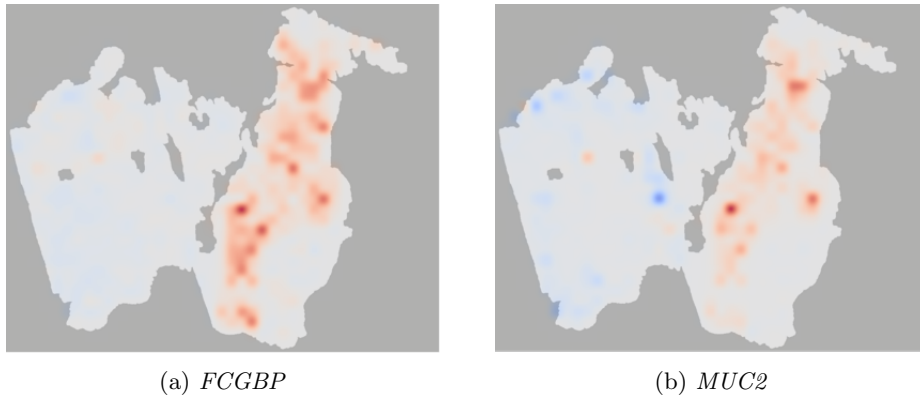


Fig. D.3: Attention-weighted predictions for co-expressed genes.