

# RETHINKING MACHINE LEARNING MODEL EVALUATION IN PATHOLOGY

Syed Ashar Javed\*, Dinkar Juyal\*, Shreya Chakraborty\*, Zahil Shanis\*, Harsha Pokkalla, Aaditya Prakash

PathAI Inc., Boston  
{firstname.lastname}@pathai.com

## ABSTRACT

Machine Learning has been applied to pathology images in research and clinical practice with promising outcomes. However, standard ML models often lack the rigorous evaluation required for clinical decisions. Machine learning techniques for natural images are ill-equipped to deal with pathology images that are significantly large and noisy, require expensive labeling, are hard to interpret, and are susceptible to spurious correlations. We propose a set of practical guidelines for ML evaluation in pathology that address the above concerns. The paper includes measures for setting up the evaluation framework, effectively dealing with variability in labels, and a recommended suite of tests to address issues related to domain shift, robustness, and confounding variables. We hope that the proposed framework will bridge the gap between ML researchers and domain experts, leading to wider adoption of ML techniques in pathology and improving patient outcomes.

## 1 INTRODUCTION

Pathology is often treated as a ground-truth for many serious diseases and a pathologist’s diagnosis is critical for a wide variety of tasks, including drug development and clinical diagnostics. The typical workflows of a pathologist are complicated, subjective and limited by what humans can evaluate. The use of machine learning in pathology has begun to transform the field with applications like clinical decision support tools (Campanella et al., 2019), improving prognostic utility over manual reads (Taylor-Weiner et al., 2021), automating slide scoring in clinical trials (Glass et al., 2021), novel biomarker discovery (Echle et al., 2021), and understanding tumor pathology (Jiang et al., 2020; Diao et al., 2021).

These models require thorough validation and verification since they directly impact patient lives and their evaluation requires a careful look at the setting in which they will be deployed. Pathology images also differ from natural images in various ways - they are massive in resolution (billions of pixels) with each magnification-level containing markedly different information (Komura & Ishikawa, 2018), have different equivariances from natural images (Veeling et al., 2018) and contain stain variations which can confound models (Yagi, 2011; Madabhushi & Lee, 2016; Tellez et al., 2019). In addition to that, exhaustive labeling of these images is prohibitive as they require expert annotations and are time-consuming. Finally, building ML models for these problems requires understanding of biological causal structures to deal with spurious features (Castro et al., 2020). These differences necessitate model evaluation guidelines specific to pathology problems to prevent life threatening model failures.

In machine learning, dataset creation, evaluation setup, choice of metrics, and the data splits used for training, validation and testing have been areas of study in themselves, and the research community depends on agreed standards for evaluating novel methods and their utility in the real world. Previous work in medical imaging has highlighted problems in transferring these evaluation methodologies, making models brittle for deployment in real-world setting (Park & Han, 2018; England & Cheng, 2019; Varoquaux & Cheplygina, 2021; Reinke et al., 2021). In this work, we propose a guideline on ML model evaluation for pathology which go beyond standard ML metrics. We define how to setup

\* equal contributions

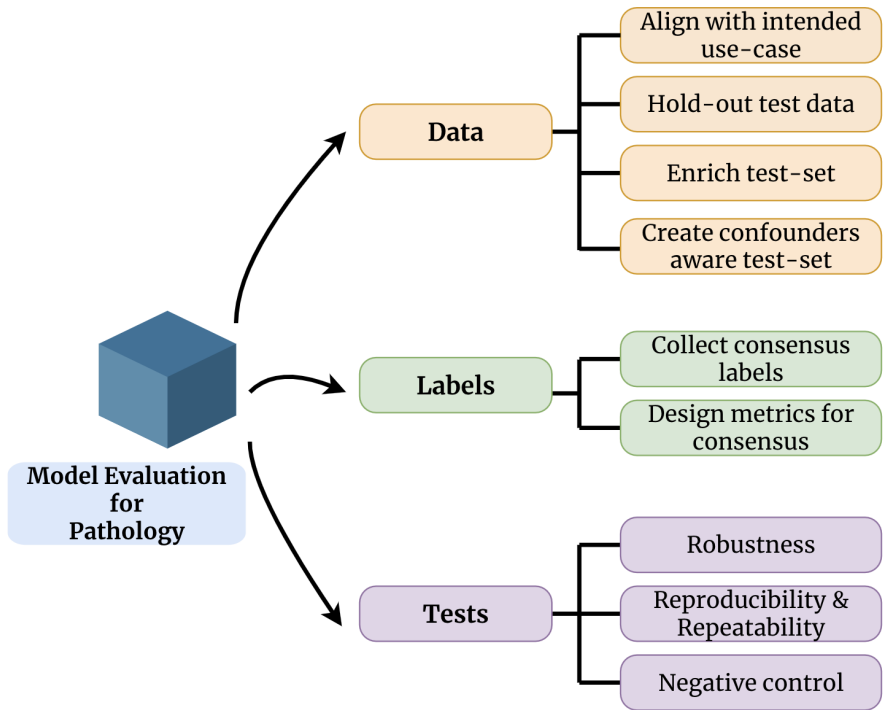


Figure 1: Guidelines for Model Evaluation

the evaluation framework, create the test-set, acquire labels for evaluation and finally suggest a set of evaluation metrics which are aimed towards real-world deployment.

## 2 GUIDELINES FOR MODEL EVALUATION

We divide model evaluation into three stages: evaluation setup, label collection, and evaluation metrics as shown in Figure 1. We provide pathology-specific recommendations for each of them in the following sections.

### 2.1 EVALUATION SETUP

#### 2.1.1 ALIGNMENT OF EVALUATION EXPERIMENTS WITH MODEL’S INTENDED USE

The applications of ML in pathology vary by their intended use, such as drug discovery, diagnostics, prognostics, decision-support, or triage. Therefore, their evaluation setup and metrics must account for the targeted use-case. For example, a clinical decision support tool is used by pathologists to look at predictions generated by the model to refine their diagnosis. The evaluation setup must measure the change in pathologist performance when assisted by the model instead of comparing model predictions against that of the pathologist.

Another example is the use of ML in building a triage tool for screening patients. In oncology, ML can be used to filter out obviously benign cases so that the pathologists prioritize malignant cases, thus improving patient care. In such scenarios, instead of using common metrics like accuracy or AUROC, one should use Precision@K%Recall due to the low tolerance for a false negative.

#### 2.1.2 ENRICHED AND CONFOUNDER-AWARE TEST SET CREATION

Pathology samples come with biological and other metadata like patient-related (age, gender, ethnicity), image-related (scanner type, image format, magnification), specimen & stain related (tissue collection type, thickness of specimen, staining protocol), and clinical (disease stage, prescribed therapy, omics data). Low prevalence of certain metadata values coupled with small, imbalanced datasets in pathology often lead to overoptimistic results if the test set composition does

not account for smaller substrata. This hidden stratification problem (Oakden-Rayner et al., 2020) can result in the model not working well for patients belonging to a hidden substrata, which in turn deteriorates trust in such tools. Another stratum which needs more representation is samples close to the decision boundary where significant ambiguity in labels exist. This occurs since labels in pathology represent a biological process which is continuous. For example, along with benign and malignant there exist borderline cases which might be under-represented if real-world data distribution is replicated in test dataset.

A related problem is the presence of confounders in the test set which are spuriously predictive of the label. For example, in pancreatic cancer, (Kather et al., 2020) show that KRAS mutation can be predicted from a pathology image. However, KRAS mutation is almost exclusively present in only one sub-type of pancreatic cancer (pancreatic ductal adenocarcinoma), thus becoming a strong biological confounder (Waters & Der, 2018). Therefore, even a well performing mutation prediction model might be relying on pancreatic sub-type features instead of the KRAS mutation features. The fact that these confounders can also be partial, non-biological and specific to the dataset, further complicates matters (Badgeley et al., 2019; Larrazabal et al., 2020; Zech et al., 2018).

To ensure that the test-set is reliable and the model is predicting the variable of interest and not a confounder, we offer the following suggestions:


- Curate test sets such that they have enough representation of any sub-strata of data that we care about or that could be a confounder (Seyyed-Kalantari et al., 2020). This can be done through enrichment (oversampling) of rare substrata as mentioned in this FDA guide (Food et al.). However practitioners should be aware of selection bias due to test-set enrichment (Yu & Eng, 2020).
- Incorporate stratified evaluation metrics as opposed to a single test-set wide global metric. For example, ensuring consistent performance across patient demographics or disease severity.
- Conduct a correlation check on all metadata values present in the dataset to see if any of them correlate highly with label values for the task of interest. For example, if all images positive for prostate cancer come from the same hospital, the model can learn to predict the hospital instead of the cancer, making it a potential confounder.

### 2.1.3 EXTERNAL HOLD-OUT TEST SET


Curated datasets are useful for research, but often suffer from the problem of a homogeneous test-set. The approach of collecting data from a single source and splitting it into train-val-test sets results in inflated performance since the models might not generalize to distributional shifts (like images from a new hospital) (Tellez et al., 2019; Campanella et al., 2019). Many popular pathology datasets (Borkowski et al., 2019; Hosseini et al., 2019; Karimi et al., 2019) and studies (Shao et al., 2021; Sudharshan et al., 2019) do not have externally held-out test sets from a new site which prevent an evaluation of the generalizability. We suggest using multiple held-out patient cohorts from different medical sites and demographics while constructing the test set (Koh et al., 2021). We also suggest checking for test data leakage in the form of patient samples being present across multiple splits. Finally, we urge the community to make such a testing setup the norm for model evaluation in pathology.

Figure illustrates the discordance among pathologists while grading NASH, an aggressive form of non-alcoholic fatty liver disease increasingly being seen as an epidemic in the US. In this case, NASH severity is found by performing a liver biopsy, viewing the extracted tissue through a microscope and observing signs of disease progression such as presence of fat, different kinds of inflammation etc.

What is most telling in this figure is the extremely low agreement among doctors while predicting disease progression even after using the most reliable way to do so, which is viewing the tissue extracted from the biopsy. The numbers on the left column denote the Kappa score for different morphological indicators of NASH (N denotes the number of patients), when the first study was done in 2005. The right column shows the agreement scores for the same from a 2020 study; there has been little or no change in these agreement scores among human experts even after 15 years.



(a) Figure illustrates the discordance among human experts while grading NASH from pathology images. NASH is an aggressive form of non-alcoholic fatty liver disease increasingly being seen as an epidemic in the US. Kappa score (agreement score) among pathologists for different morphological indicators of NASH are shown 15 years apart; there has been no improvement in pathologist consensus.



(b) Variations in pathology images across scanners and lab sites. The first row shows the same slide scanned with different scanners with visually perceptible differences. The second row shows slides collected from the same scanner but different lab sites, which differ visually due to staining protocols, choice of reagents, method of specimen preparation and so on. Models struggle to generalize to these variations

Figure 2

## 2.2 EVALUATION LABELS

It is possible to get fast and cheap labels using crowdsourced tools for natural images, and the correctness of these labels is easily verifiable. In contrast, it is a lot harder to collect ground-truth labels for pathology. Pathology images have billions of pixels, millions of cells, and thousands of regions of interest, making label collection extremely resource-intensive. Annotators (typically certified pathologists) are in short supply and per unit cost of annotations is order of magnitude higher than natural images. Annotators also have their own biases (Aeffner et al., 2017).

A more insidious problem is of the high ‘inter-annotator variability’ (Carrasco-Zevallos et al., 2021) as shown in Figure 2a, and ‘intra-annotator variability’ (Kleiner et al., 2005; Davison et al., 2020): annotators disagree not only with each other but also with their past selves. These variabilities exist at different stages of label generation - they can originate from inherent ambiguity in the labeling system, different interpretations of the same labeling scheme by different experts, and the differences in the expertise of these annotators (Ozkan et al., 2016). The threshold for acceptable variability is subjective, and depends on factors such as the particular organ and disease.

We propose the following guidelines while collecting labels:

- Collect multiple annotations per image and measure the inter-annotator agreement. If possible, collect multiple annotations on the same image from the same annotator at adequately-spaced timepoints (called ‘washout period’) to measure intra-annotator variability. If these variabilities are high, the ML task might not be viable.
- Aggregate the labels to a single consensus if the intra and inter-annotator variabilities are low. Instead of comparing model performance with a single pathologist (Liu et al., 2017) or consensus, measure the agreement of the model with the consensus and judge it against inter-annotator consensus. This is more realistic and appreciates the difficulty of the task as judged by human’s inability to do it consistently (Raciti et al., 2020). Kappa (Linear, Quadratic) and Intraclass Correlation Coefficient (ICC) are popularly used to measure single and multi-rater agreement Bulten et al. (2022); Koo & Li (2016).

## 2.3 EVALUATION METRICS

### 2.3.1 REPRODUCIBILITY, REPEATABILITY, & ROBUSTNESS TESTS

Pathology images are heterogeneous with regard to staining, tissue thickness, and are prone to variations during tissue processing, cutting, staining, and digitization (Figure 2b). These variations cause minor visually perceptible changes to the image, but can drastically change the model prediction. Lack of robustness to these perturbations poses a major challenge to the wider adoption

and use of ML in pathology, especially in a clinical setting (Schömig-Markiefka et al., 2021). Accuracy based assessment of models, although necessary, is not sufficient for capturing model generalization. Following tests must be incorporated into the model evaluation framework to assess model’s robustness to these variations:

- **Reproducibility Tests:** Models must be tested under known sources of variations and evaluated for consistency and accuracy in performance. This includes testing by explicitly changing external variations such as staining & scanning, variations associated with hardware (non-determinism in GPU) and software packages. Models in pathology are usually trained on smaller patches of the image due to their large size (Shao et al., 2021; Taylor-Weiner et al., 2021). Stochasticity introduced due to sampling the patches can lead to inconsistent predictions and robustness to this variability must also be reported.
- **Repeatability Tests:** Repeatability expresses the precision under the same operating conditions over a short interval of time (FDA, 1995). This test captures unknown & unaccountable variations that are not covered by reproducibility tests. Consistency in model predictions must be measured and reported by rescanning & rerunning inference on the same image multiple times, keeping all known sources of variations constant.
- **Robustness to Perturbations:** Creating test sets that capture all possible real-world variations of images is not practical. An alternative can be to evaluate robustness on synthetically generated perturbations inspired by real world variations to image samples, by measuring consistency of predictions across these perturbations (Hendrycks & Dietterich, 2019; Schömig-Markiefka et al., 2021). Models that are robust to synthetic variations have been shown to be robust to natural variations (Faryna et al., 2021).

### 2.3.2 NEGATIVE CONTROL TESTS

Cross domain differences in pathology images collected from different sources may confound the discovery of explanatory variables and lead to poor generalization. This phenomenon, called ‘batch effects’ (Goh et al., 2017), can lead to the model learning spurious features. Existing literature on batch effects show that models can predict patient age, sex, race, time of slide preparation, site, and scanner (Schmitt et al., 2021; Howard et al., 2020).

Although creating confounder-aware test sets can mitigate some of these issues, partial and hidden confounders may go unnoticed and affect model performance. Another approach to tackle this issue is to perform negative control tests which involve running experiments under conditions in which the model is not expected to produce correct predictions (Lipsitch et al., 2010). We suggest the following strategies for designing negative control experiments for ML in pathology:

- Train with regions of the image which have no discriminative information. For example, a model trained to predict cancer using only background (non-tissue area) should perform poorly.
- Train without regions which are essential for predicting the label. For example, models for predicting cancer should perform poorly when trained and deployed on patches sampled from non-cancer regions.

Failure of a negative control test signals the presence of confounders and biases that create spurious correlations. These tests are not designed to detect all types of confounders, but they help detect the prominent ones and take appropriate actions.

## 3 CONCLUSION

Machine learning for pathology requires different evaluation standards due to differences in image characteristics and clinical use-case. Existing datasets and evaluation metrics prevalent within the ML community need to be rethought to fit pathology applications. We have observed that there exist two distinct communities, one of ML researchers focused on the development of novel methods and the other of computational pathology researchers focused on exploring and utilizing the power of ML in healthcare, and model evaluation differs between them materially. We argue that domain-specific knowledge from pathology concerning label reliability, test set enrichment, and biological confounders should permeate into ML research. Conversely, advancements in ML

such as synthetic data generation, robustness and sensitivity analysis, and out-of-domain testing deserve wider adoption in computational pathology. In this paper, we propose a set of guidelines for evaluating ML models in pathology which will help researchers create better datasets, metrics, and evaluation setup. We believe this will improve patient outcomes through ML applications in pathology and inform real-world motivated research.

#### ACKNOWLEDGMENTS

Authors would like to thank Jonathan Rubin (Philips, Amazon) and Maz Abulnaga (MIT) for their inputs.

#### REFERENCES

- Famke Aeffner, Kristin Wilson, Nathan T. Martin, Joshua C. Black, Cris L. Luengo Hendriks, Brad Bolon, Daniel G. Rudmann, Roberto Gianani, Sally R. Koegler, Joseph Krueger, and G. Dave Young. The Gold Standard Paradox in Digital Image Analysis: Manual Versus Automated Scoring as Ground Truth. *Archives of Pathology & Laboratory Medicine*, 141(9):1267–1275, 05 2017. ISSN 0003-9985. doi: 10.5858/arpa.2016-0386-RA. URL <https://doi.org/10.5858/arpa.2016-0386-RA>.
- Marcus A Badgeley, John R Zech, Luke Oakden-Rayner, Benjamin S Glicksberg, Manway Liu, William Gale, Michael V McConnell, Bethany Percha, Thomas M Snyder, and Joel T Dudley. Deep learning predicts hip fracture using confounding patient and healthcare variables. *NPJ digital medicine*, 2(1):1–10, 2019.
- Andrew A Borkowski, Marilyn M Bui, L Brannon Thomas, Catherine P Wilson, Lauren A DeLand, and Stephen M Mastorides. Lung and colon cancer histopathological image dataset (lc25000). *arXiv preprint arXiv:1912.12142*, 2019.
- Wouter Bulten, Kimmo Kartasalo, Po-Hsuan Cameron Chen, Peter Ström, Hans Pinckaers, Kunal Nagpal, Yuannan Cai, David F Steiner, Hester van Boven, Robert Vink, et al. Artificial intelligence for diagnosis and gleason grading of prostate cancer: the panda challenge. *Nature medicine*, pp. 1–10, 2022.
- Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Miraflor, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine*, 25(8):1301–1309, 2019.
- Oscar Carrasco-Zevallos, Amaro Taylor-Weiner, Harsha Pokkalla, Maryam Pouryahya, Charles Biddle-Snead, Ling Han, Ryan Huss, Dinkar Juyal, Zahil Shanis, Aryan Pedawi, et al. Ai-based histologic measurement of nash (aim-nash): A drug development tool for assessing clinical trial end points. In *JOURNAL OF HEPATOLOGY*, volume 75, pp. S254–S254. ELSEVIER RADARWEG 29, 1043 NX AMSTERDAM, NETHERLANDS, 2021.
- Daniel C Castro, Ian Walker, and Ben Glocker. Causality matters in medical imaging. *Nature Communications*, 11(1):1–10, 2020.
- Beth A. Davison, Stephen A. Harrison, Gad Cotter, Naim Alkhouri, Arun Sanyal, Christopher Edwards, Jerry R. Colca, Julie Iwashita, Gary G. Koch, and Howard C. Dittrich. Suboptimal reliability of liver biopsy evaluation has implications for randomized clinical trials. *Journal of Hepatology*, 73(6):1322–1332, 2020. ISSN 0168-8278. doi: <https://doi.org/10.1016/j.jhep.2020.06.025>. URL <https://www.sciencedirect.com/science/article/pii/S0168827820303998>.
- James A Diao, Jason K Wang, Wan Fung Chui, Victoria Mountain, Sai Chowdary Gullapally, Ramprakash Srinivasan, Richard N Mitchell, Benjamin Glass, Sara Hoffman, Sudha K Rao, et al. Human-interpretable image features derived from densely mapped cancer pathology slides predict diverse molecular phenotypes. *Nature communications*, 12(1):1–15, 2021.

- Amelie Echle, Niklas Timon Rindtorff, Titus Josef Brinker, Tom Luedde, Alexander Thomas Pearson, and Jakob Nikolas Kather. Deep learning in cancer pathology: a new generation of clinical biomarkers. *British journal of cancer*, 124(4):686–696, 2021.
- Joseph R England and Phillip M Cheng. Artificial intelligence for medical image analysis: a guide for authors and reviewers. *American journal of roentgenology*, 212(3):513–519, 2019.
- Khrystyna Faryna, Jeroen van der Laak, and Geert Litjens. Tailoring automated data augmentation to h&e-stained histopathology. In *Medical Imaging with Deep Learning*, 2021.
- FDA. Guideline for Industry: Text on Validation of Analytical Procedures. <https://www.fda.gov/media/71724/download>, 1995.
- Food, Drug Administration, et al. Guidance for industry and food and drug administration staff: Computer-assisted detection devices applied to radiology images and radiology device data-premarket notification [510 (k)] submissions.
- Benjamin Glass, Michel Erminio Vandenberghe, Surya Teja Chavali, Syed Ashar Javed, Marlon Rebelatto, Shamira Sridharan, Hunter Elliott, Sudha Rao, Michael Montalto, Murray Resnick, et al. Machine learning models to quantify her2 for real-time tissue image analysis in prospective clinical trials., 2021.
- Wilson Wen Bin Goh, Wei Wang, and Limsoon Wong. Why batch effects matter in omics data, and how to avoid them. *Trends in biotechnology*, 35(6):498–507, 2017.
- Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.
- Mahdi S Hosseini, Lyndon Chan, Gabriel Tse, Michael Tang, Jun Deng, Sajad Norouzi, Corwyn Rowsell, Konstantinos N Plataniotis, and Savvas Damaskinos. Atlas of digital pathology: A generalized hierarchical histological tissue type-annotated database for deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11747–11756, 2019.
- Frederick M Howard, James Dolezal, Sara Kochanny, Jefree Schulte, Heather Chen, Lara Heij, Dezheng Huo, Rita Nanda, Olufunmilayo I Olopade, Jakob N Kather, et al. The impact of digital histopathology batch effect on deep learning model accuracy and bias. *bioRxiv*, 2020.
- Yahui Jiang, Meng Yang, Shuhao Wang, Xiangchun Li, and Yan Sun. Emerging role of deep learning-based artificial intelligence in tumor pathology. *Cancer communications*, 40(4):154–166, 2020.
- Davood Karimi, Guy Nir, Ladan Fazli, Peter C Black, Larry Goldenberg, and Septimiu E Salcudean. Deep learning-based gleason grading of prostate cancer from histopathology images—role of multiscale decision aggregation and data augmentation. *IEEE journal of biomedical and health informatics*, 24(5):1413–1426, 2019.
- Jakob Nikolas Kather, Lara R Heij, Heike I Grabsch, Chiara Loeffler, Amelie Echle, Hannah Sophie Muti, Jeremias Krause, Jan M Niehues, Kai AJ Sommer, Peter Bankhead, et al. Pan-cancer image-based detection of clinically actionable genetic alterations. *Nature Cancer*, 1(8):789–799, 2020.
- David E. Kleiner, Elizabeth M. Brunt, Mark Van Natta, Cynthia Behling, Melissa J. Contos, Oscar W. Cummings, Linda D. Ferrell, Yao-Chang Liu, Michael S. Torbenson, Aynur Unalp-Arida, Matthew Yeh, Arthur J. McCullough, Arun J. Sanyal, and Nonalcoholic Steatohepatitis Clinical Research Network. Design and validation of a histological scoring system for nonalcoholic fatty liver disease. *Hepatology*, 41(6):1313–1321, 2005. doi: <https://doi.org/10.1002/hep.20701>. URL <https://aasldpubs.onlinelibrary.wiley.com/doi/abs/10.1002/hep.20701>.
- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021.

- Daisuke Komura and Shumpei Ishikawa. Machine learning methods for histopathological image analysis. *Computational and structural biotechnology journal*, 16:34–42, 2018.
- Terry K Koo and Mae Y Li. A guideline of selecting and reporting intraclass correlation coefficients for reliability research. *Journal of chiropractic medicine*, 15(2):155–163, 2016.
- Agostina J Larrazabal, Nicolás Nieto, Victoria Peterson, Diego H Milone, and Enzo Ferrante. Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences*, 117(23):12592–12594, 2020.
- Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology (Cambridge, Mass.)*, 21(3):383, 2010.
- Yun Liu, Krishna Gadepalli, Mohammad Norouzi, George E. Dahl, Timo Kohlberger, Aleksey Boyko, Subhashini Venugopalan, Aleksei Timofeev, Philip Q. Nelson, Greg S. Corrado, Jason D. Hipp, Lily Peng, and Martin C. Stumpe. Detecting cancer metastases on gigapixel pathology images, 2017.
- Anant Madabhushi and George Lee. Image analysis and machine learning in digital pathology: Challenges and opportunities. *Medical image analysis*, 33:170–175, 2016.
- Luke Oakden-Rayner, Jared Dunnmon, Gustavo Carneiro, and Christopher Ré. Hidden stratification causes clinically meaningful failures in machine learning for medical imaging. In *Proceedings of the ACM conference on health, inference, and learning*, pp. 151–159, 2020.
- Tayyar A Ozkan, Ahmet T Eruyar, Oguz O Cebeci, Omur Memik, Levent Ozcan, and Ibrahim Kuskonmaz. Interobserver variability in gleason histological grading of prostate cancer. *Scandinavian journal of urology*, 50(6):420–424, 2016.
- Seong Ho Park and Kyunghwa Han. Methodologic guide for evaluating clinical performance and effect of artificial intelligence technology for medical diagnosis and prediction. *Radiology*, 286(3):800–809, 2018.
- Patricia Raciti, Jillian Sue, Rodrigo Ceballos, Ran Godrich, Jeremy D Kunz, Supriya Kapur, Victor Reuter, Leo Grady, Christopher Kanan, David S Klimstra, et al. Novel artificial intelligence system increases the detection of prostate cancer in whole slide images of core needle biopsies. *Modern Pathology*, 33(10):2058–2066, 2020.
- Annika Reinke, Matthias Eisenmann, Minu D Tizabi, Carole H Sudre, Tim Rädtsch, Michela Antonelli, Tal Arbel, Spyridon Bakas, M Jorge Cardoso, Veronika Cheplygina, et al. Common limitations of image processing metrics: A picture story. *arXiv preprint arXiv:2104.05642*, 2021.
- Max Schmitt, Roman Christoph Maron, Achim Hekler, Albrecht Stenzinger, Axel Hauschild, Michael Weichenthal, Markus Tiemann, Dieter Krahl, Heinz Kutzner, Jochen Sven Utikal, et al. Hidden variables in deep learning digital pathology and their potential to cause batch effects: prediction model study. *Journal of medical Internet research*, 23(2):e23436, 2021.
- Birgid Schömig-Markiefka, Alexey Pryalukhin, Wolfgang Hulla, Andrey Bychkov, Junya Fukuoka, Anant Madabhushi, Viktor Achter, Lech Nieroda, Reinhard Büttner, Alexander Quaas, et al. Quality control stress test for deep learning-based diagnostic model in digital pathology. *Modern Pathology*, 34(12):2098–2108, 2021.
- Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew McDermott, Irene Y Chen, and Marzyeh Ghassemi. Chexclusion: Fairness gaps in deep chest x-ray classifiers. In *BIOCOMPUTING 2021: Proceedings of the Pacific Symposium*, pp. 232–243. World Scientific, 2020.
- Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, et al. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *Advances in Neural Information Processing Systems*, 34, 2021.
- PJ Sudharshan, Caroline Petitjean, Fabio Spanhol, Luiz Eduardo Oliveira, Laurent Heutte, and Paul Honeine. Multiple instance learning for histopathological breast cancer image classification. *Expert Systems with Applications*, 117:103–111, 2019.



- Amaro Taylor-Weiner, Harsha Pokkalla, Ling Han, Catherine Jia, Ryan Huss, Chuhan Chung, Hunter Elliott, Benjamin Glass, Kishalve Pethia, Oscar Carrasco-Zevallos, et al. A machine learning approach enables quantitative measurement of liver histology and disease monitoring in nash. *Hepatology*, 74(1):133–147, 2021.
- David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen Van Der Laak. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical image analysis*, 58:101544, 2019.
- Gaël Varoquaux and Veronika Cheplygina. How i failed machine learning in medical imaging—shortcomings and recommendations. *arXiv preprint arXiv:2103.10292*, 2021.
- Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant cnns for digital pathology. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 210–218. Springer, 2018.
- Andrew M Waters and Channing J Der. Kras: the critical driver and therapeutic target for pancreatic cancer. *Cold Spring Harbor perspectives in medicine*, 8(9):a031435, 2018.
- Yukako Yagi. Color standardization and optimization in whole slide imaging. In *Diagnostic pathology*, volume 6, pp. 1–12. Springer, 2011.
- Alice C Yu and John Eng. One algorithm may not fit all: how selection bias affects machine learning performance. *Radiographics*, 40(7):1932–1937, 2020.
- John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric Karl Oermann. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS medicine*, 15(11):e1002683, 2018.