

Article

xDEEP-MSI: Explainable Bias-Rejecting Microsatellite Instability Deep Learning System in Colorectal Cancer

Aurelia Bustos ^{1,*} , Artemio Payá ^{2,3}, Andrés Torrubia ¹, Rodrigo Jover ^{2,3}, Xavier Llor ⁴, Xavier Bessa ⁵, Antoni Castells ⁶, Ángel Carracedo ⁷  and Cristina Alenda ^{2,3,*} 

- ¹ AI Cancer Research Unit, MedBravo, 03560 Alicante, Spain; andres@medbravo.org
² Pathology Department, Alicante University General Hospital (HGUA), 03010 Alicante, Spain; paya_art@gva.es (A.P.); jover_rod@gva.es (R.J.)
³ Alicante Institute for Health and Biomedical Research (ISABIAL), 03010 Alicante, Spain
⁴ Department of Medicine, Yale Cancer Center, Yale University, New Haven, CT 06511, USA; xavier.llor@yale.edu
⁵ Gastroenterology Department, Hospital del Mar Medical Research Institute (IMIM), 08003 Barcelona, Spain; xbessa@parcdesalutmar.cat
⁶ Centro de Investigación Biomédica en Red de Enfermedades Hepáticas y Digestivas (CIBERehd), Gastroenterology Department, Institut d'Investigacions Biomèdiques August Pi i Sunyer (IDIBAPS), Hospital Clinic, University of Barcelona, 08036 Barcelona, Spain; CASTELLS@clinic.cat
⁷ Fundación Pública Gallega de Medicina Genómica (GMX), 15706 Santiago de Compostela, Spain; angel.carracedo@usc.es
* Correspondence: aurelia@medbravo.org (A.B.); alenda_cri@gva.es (C.A.)



check for updates

Citation: Bustos, A.; Payá, A.; Torrubia, A.; Jover, R.; Llor, X.; Bessa, X.; Castells, A.; Carracedo, Á.; Alenda, C. xDEEP-MSI: Explainable Bias-Rejecting Microsatellite Instability Deep Learning System in Colorectal Cancer. *Biomolecules* **2021**, *11*, 1786. <https://doi.org/10.3390/biom11121786>

Academic Editor: William Cho

Received: 3 November 2021

Accepted: 26 November 2021

Published: 29 November 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract: The prediction of microsatellite instability (MSI) using deep learning (DL) techniques could have significant benefits, including reducing cost and increasing MSI testing of colorectal cancer (CRC) patients. Nonetheless, batch effects or systematic biases are not well characterized in digital histology models and lead to overoptimistic estimates of model performance. Methods to not only palliate but to directly abrogate biases are needed. We present a multiple bias rejecting DL system based on adversarial networks for the prediction of MSI in CRC from tissue microarrays (TMAs), trained and validated in 1788 patients from EPICOLON and HGUA. The system consists of an end-to-end image preprocessing module that tile samples at multiple magnifications and a tissue classification module linked to the bias-rejecting MSI predictor. We detected three biases associated with the learned representations of a baseline model: the project of origin of samples, the patient's spot and the TMA glass where each spot was placed. The system was trained to directly avoid learning the batch effects of those variables. The learned features from the bias-ablated model achieved maximum discriminative power with respect to the task and minimal statistical mean dependence with the biases. The impact of different magnifications, types of tissues and the model performance at tile vs patient level is analyzed. The AUC at tile level, and including all three selected tissues (tumor epithelium, mucin and lymphocytic regions) and 4 magnifications, was 0.87 ± 0.03 and increased to 0.9 ± 0.03 at patient level. To the best of our knowledge, this is the first work that incorporates a multiple bias ablation technique at the DL architecture in digital pathology, and the first using TMAs for the MSI prediction task.

Keywords: digital pathology; deep neural networks; bias ablation; adversarial networks; colorectal carcinoma; microsatellite instability

1. Introduction

Approximately 3% of colorectal cancers (CRC) arise in the context of Lynch syndrome (LS), where the patient has a germline mutation in a DNA mismatch repair (MMR) gene [1]. Historically, CRC patients were tested for LS if they were at high risk according to clinical criteria, e.g., aged under 50 years or with a strong family history. Several clinicopathologic criteria (e.g., Amsterdam criteria, revised Bethesda guidelines) were used to identify

individuals at risk for Lynch syndrome or eligible for tumor-based MSI testing [2]. However, a large proportion of LS patients were missed by this strategy [3]. Currently, new diagnostics guidance are recommending that all patients with newly diagnosed CRC be screened for LS. Universal tumor-based genetic screening for Lynch syndrome, with MSI screening by PCR or defective MMR(dMMR) detection by IHC testing of all CRCs regardless of age, has greater sensitivity for identification of Lynch syndrome as compared with other strategies [1]. The pathway includes testing tumour tissue for defective MMR by either microsatellite instability (MSI) testing or immunohistochemistry (IHC) for the MMR proteins MLH1, PMS2, MSH2 and MSH6. Tumours showing MSI or MLH1 loss should subsequently undergo *BRAF* mutation testing followed by *MLH1* promoter methylation analysis in the absence of a *BRAF* mutation. Patients with tumours showing MSH2, MSH6 or isolated PMS2 loss, or MLH1 loss/MSI with no evidence of *BRAF* mutation/*MLH1* promoter hypermethylation, are referred for germline testing if clinically appropriate.

MSI is not specific for LS, and approximately 15 percent of all sporadic CRCs and 5 to 10 percent of metastatic CRCs demonstrate MSI due to hypermethylation of *MLH1* [3–5]. Sporadic microsatellite instability-high (MSI-H) CRCs typically develop through a methylation pathway called CpG island methylator phenotype (CIMP), which is characterized by aberrant patterns of DNA methylation and frequently by mutations in the *BRAF* gene. These cancers develop somatic promoter methylation of *MLH1*, leading to loss of MLH1 function and resultant MSI. The prevalence of loss of MLH1 expression in CRC increases markedly with aging and this trend is particularly evident in women [6].

While guidelines set forth by multiple professional societies recommend universal testing for dMMR/MSI [7], these methods require additional resources and are not available at all medical facilities, so many CRC patients are not currently tested [8].

Since the last two decades, certain histology-based prediction models that rely on hand-crafted clinico-pathologic feature extraction—such as age < 50, female sex, right sided location, size \geq 60 mm, *BRAF* mutation, tumor infiltrating lymphocytes (TILs), a peritumoral lymphocytic reaction, mucinous morphology and increased stromal plasma cells—have reported encouraging performance but has not been sufficient to supersede universal testing for MSI/dMMR [9]. Measurement of the variables for MSI prediction, requires significant effort and expertise by pathologists, and inter-rater differences may affect the perceived reliability of histology-based scoring systems [9]. However, this work is fundamental to the premise that MSI can be predicted from histology, which was recently proposed as a task for deep learning from digital pathology [10,11].

Research on deep learning methods to predict MSI directly from hematoxylin and eosin (H&E) stained slides of CRC have proliferated in the last years, refs. [10–17] with reported accuracy rapidly improved on most recent works [14,15].

If successful, this approach could have significant benefits, including reducing cost and resource-utilization and increasing the proportion of CRC patients that are tested for MSI. Additional potential benefits are to increase the capability of detecting MSI over current methods alone. Some tumors are either dMMR or MSI-H but not the other. Testing of tumors with only immunohistochemistry (IHC) or polymerase chain reaction (PCR) will falsely exclude some patients from immunotherapy [18]. A system trained on both techniques could overcome this limitation, obviating co-testing with both MMR IHC and MSI PCR as a screening strategy for evaluating the eligibility status for immunotherapy.

Among current limitations, the developed systems so far are not able to distinguish between somatic and germline etiology of MSI, such that confirmatory testing is required. Another important limitation is generalizability due to batch effects, while those systems have proved excellent performance on well curated cohorts that are similar to training data, the performance is not robust to differing patient and tissue characteristics. This limitation is evident from the performance deterioration when systems trained on a single datasource are tested on external datasets [14]. This limitation can be palliated by using larger multi-institutional datasets from different institutions for training, as shown by [15] that estimated 5000 to be the optimal number of patients needed for this specific task.

Nonetheless, compiling such international large scale datasets is costly and unfeasible in most cases and does not eliminate batch effects. Regardless of the dataset size and number of contributing sites, the propensity for overfitting of digital histology models to site level characteristics is incompletely characterized and is infrequently accounted for in internal validation of deep learning models [19]. For example assessments of stain normalization and augmentation techniques have focused on the performance of models in validation sets, rather than true elimination of batch effect [15,20]. In addition to staining techniques, batch effects originate from other multiple reasons, such as digitization of slides, variations due to scanner calibration and choice of resolution and magnification. Batch effects in training, validation and testing, must be accounted for to ensure equitable application of DL. Batch effects leads to overoptimistic estimates of model performance and methods to not only palliate but to directly abrogate this bias are needed [19].

Herein, we present a novel approach in digital pathology to eliminate the batch effects at the deep learning architecture following the methodology described in [21], where by means of an adversarial training and bias distillation regime, the model avoids learning undesirable characteristics of datasets such as the hospital or other protected variables. Adversarial training has been previously explored to improve the generalizability of predictive model to predict MSI status on tumor types not observed in training [10], nonetheless this technique has not yet been applied for removing multiple batch effects in digital pathology. We extend the methodology described in [21] by systematically assessing and quantifying the spurious associations of protected variables (biases) on the network and then leveraging a multiple bias-ablation architecture in the model.

The remainder of the paper is as follows. Section 2 describes the methodology employed, including the study population, the image preprocessing module, partitioning methods and deep learning architecture, the identification of biases, the bias-ablation system and training regime. Section 3 first describes the results of the tissue classifier module, the image dataset obtained after preprocessing, the results of the bias identification, the results of the MSI-status classifier module at image level with the demonstration of batch effect distillation, performance results at patient level, analysing the impact of types of tissues included and image magnifications, and the explainability of predictions. Finally, Section 4 addresses the discussion, conclusions and future work.

2. Material and Methods

2.1. Study Population, Data Collection and Ground Truth Ascertainment

The study population initially consisted of the H&E of 57 tissue-microarrays (TMAs) that included cylindrical tissue samples of 1 mm diameter each (spots), in duplicate from all patients prospectively collected in the EPICOLON project. TMAs were constructed from tumor paraffin blocks from each hospital. For the present project, the paraffin blocks from each hospital were sectioned and the same hematoxylin eosin staining was performed on all the cases included. In this technique, each TMA glass consists of multiple spots placed in one microscope glass slide. The same glass slide vendor was used for all the TMAs. The EPICOLON project was a population-based, observational, cohort study which included 1705 patients with CRC from 2 Spanish nationwide multi-center studies: EPICOLON I and EPICOLON II. EPICOLON I included consecutive patients with a new diagnosis of CRC between November 2000 and October 2001 with the main goal of estimating the incidence of LS in Spain [22]. EPICOLON II also included consecutive patients with newly diagnosed CRC between March 2006 and December 2007 and from 2009 to 2013 with the aim of investigating different aspects related to the diagnosis of hereditary CRC [23]. For both cohorts, the inclusion criteria were all patients with a de-novo histologically confirmed diagnosis of colorectal adenocarcinoma and who attended 25 teaching and community hospitals across Spain during the different recruiting periods covered by EPICOLON I and EPICOLON II as previously detailed. Exclusion criteria for both studies were patients in whom CRC developed in the context of familial adenomatous polyposis or inflammatory bowel disease, and patient or family refusal to participate in the study. It was assumed

that the EPICOLON population was representative of the Spanish population, due to the large number of participating centres (most of them referral centres of each area), their homogeneous distribution throughout the country, and the lack of ethnic differences among regions [22]. Both studies were approved by the institutional review boards of the participating hospitals. The overall MSI frequency in the EPICOLON project was 7.4%.

The population was further expanded with 283 additional patients retrospectively obtained from the *Hospital Universitario de Alicante, Spain* (HGUA) from 2017 onwards, which -once preprocessed- added 66 MSI-H and 177 microsatellite stable (MSS) cases to the final study population as shown in Figure 1. The demographic and main clinical characteristics of the final study sample after preprocessing are summarized in Table 1. The corresponding H&E images were provided in 15 TMAs that included 1 mm spots in duplicate for each patient.

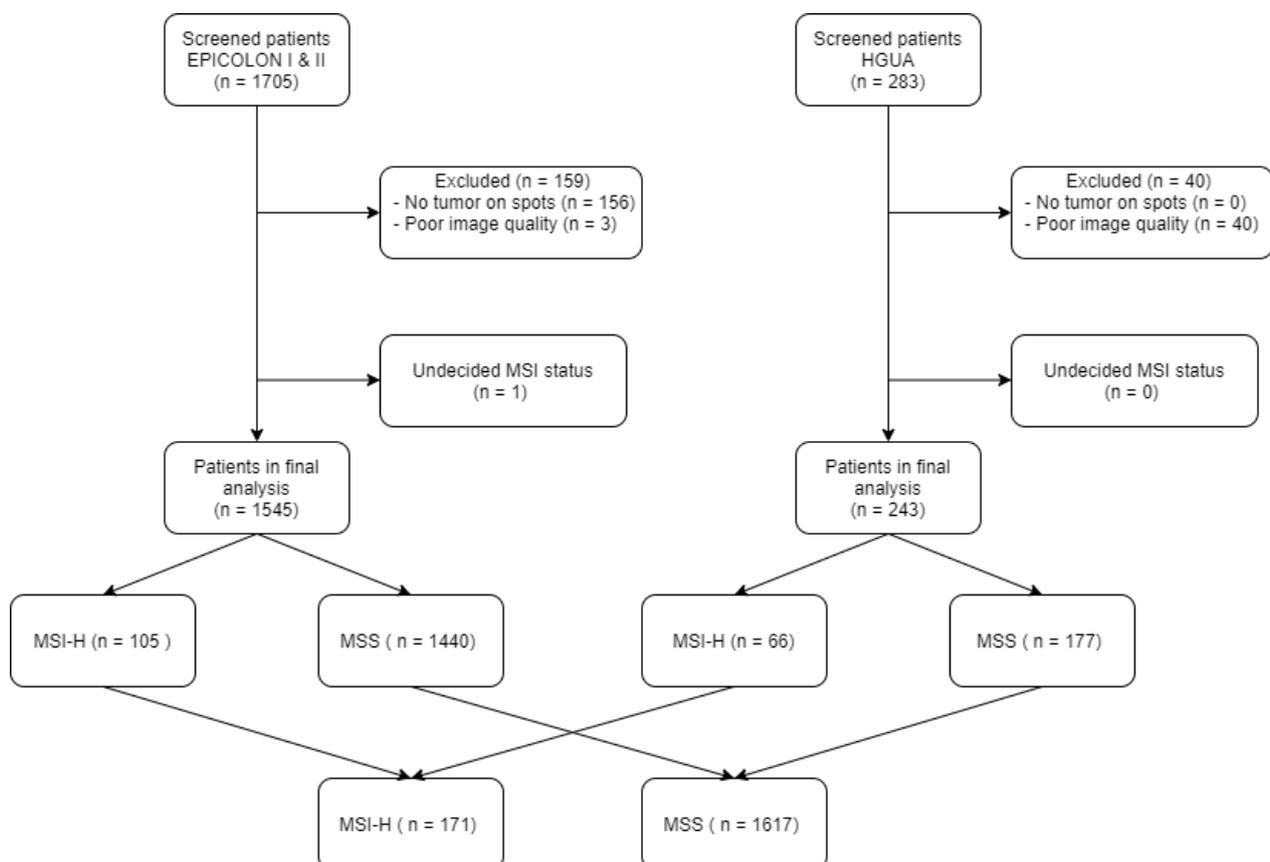


Figure 1. Study sample flowchart: For image preprocessing, the tissue classifier module was used to exclude images based on the criteria of no remaining tumor epithelium on spots.

For the task of MSI prediction each patient and corresponding spots were labeled as MSI-H vs. MSS. MSI-H was defined as tumour-tissue testing defective MMR(dMMR) by either microsatellite instability (MSI) testing or immunohistochemistry (IHC) for the MMR proteins MLH1, PMS2, MSH2 or MSH6. Tumours showing MSI or MLH1 loss underwent *BRAF* mutation testing followed by *MLH1* promoter methylation analysis in the absence of a *BRAF* mutation. Briefly, patients with tumours showing MSH2, MSH6 or isolated PMS2 loss, or MLH1 loss/MSI with no evidence of *BRAF* mutation (regardless of presence or not of *MLH1* promoter hypermethylation), were labeled as MSI-H and all others as MSS.

TMAs were scanned with the VENTANA ROCHE iSCAN scanner at magnification $\times 40$ corresponding to a maximal resolution of 0.25 microns per pixel (MPP).

Table 1. Patient Characteristics (n = 1788): Main demographic and clinical characteristics of patients included in the final study sample after image preprocessing. * Dukes' stage was not available for the whole study sample.

Characteristic	n
Age (years)	75.6 (27–101)
Gender	
Male	1037 (58.1%)
Female	749 (41.9%)
Personal history of cancer	177 (9.8%)
Dukes' Stage *	
I	119 (6.6%)
II	370 (20.6%)
III	271 (15.1%)
IV	131 (7.3%)

This project was approved by the institutional research committee CEIm PI2019-029 from ISABIAL and both the images and associated clinical information was previously anonymized. The data from EPICOLON project and HGUA are not publicly available, in accordance with the research group and institutional requirements governing human subject privacy protection.

The NCT-CRC-HE-100K and CRC-VAL-HE-7K datasets, consisting of 100,000 non-overlapping image patches from hematoxylin & eosin (H&E) stained histological images of human colorectal cancer (CRC) and normal tissue labeled with 9 tissue classes—adipose (ADI), background without tissue (BACK), debris (DEB), lymphocytes (LYM), mucus (MUC), smooth muscle (MUS), normal colon epithelium (NORM), cancer-associated stroma (STR) and colorectal adenocarcinoma epithelium (TUM)—were used to train a model in the task of tissue class prediction. All images were 224×224 pixels (px) at 0.5 MPPs. All images were color-normalized using Mazenko's method [24]. All image tiles for the NCT-CRC-HE-100K and CRC-VAL-HE-7K datasets are available online at <https://zenodo.org/record/1214456#.XcNpCpNKjyw> (accessed on 25 October 2019).

2.2. Preprocessing Module

The image preprocessing module consisted of a TMA-customized dynamic extractor of tiles of 400×400 pixels corresponding to adjacent regions without overlap parameterized at different magnifications ($\times 40$, $\times 20$, $\times 10$, $\times 5$, $\times 0$) linked to a DL tissue classifier (described in see Section 2.3) that filtered spots without viable colorectal adenocarcinoma epithelium and at the same time selected the regions of interest based on type of tissues.

2.3. Deep Learning Architecture

The DL architecture (see Figure 2) consisted of an end-to-end deep learning system linking three types of modules: (1) A classifier of 9 types of tissues at tile level: adipose (ADI), background without tissue (BACK), debris (DEB), lymphocytes (LYM), mucus (MUC), smooth muscle (MUS), normal colon epithelium (NORM), cancer-associated stroma (STR) and colorectal adenocarcinoma epithelium (TUM) (2) the MSI status classifier *MSI* and (3) the batch effect module learners *BE*, which has as many learners or heads as required to account for multiple biases.

The classifier of tissues, had as a convolutional backbone a Resnet34 pre-trained on Imagenet, and two hidden layer of dimension 512 with ReLU as the activation function, and 9 out features. This module was trained with the NCT-CRC-HE-100K and CRC-VAL-HE-7K datasets on the final task of learning the 9 types of tissue. Once trained, it was used both to select the regions of interest (ROIs) based on type of tissues in the preprocessing

module and as the feature extractor backbone FE in the end-to-end system (removing the head) connected with the MSI and BE modules. The FE yielded 512 intermediate features.

Both the MSI and all BE modules have two hidden layers of dimension 512 with ReLU as the activation function, and 2 out features.

The DL system without the BE modules is referred herein as the baseline model, and when including the BE modules, is referred as the bias-ablated model.

All code was programmed using the Python programming language. Machine and deep learning methods were implemented using FastAI [25] and PyTorch. QPath v0.2.3 [26] was used for annotation purposes.

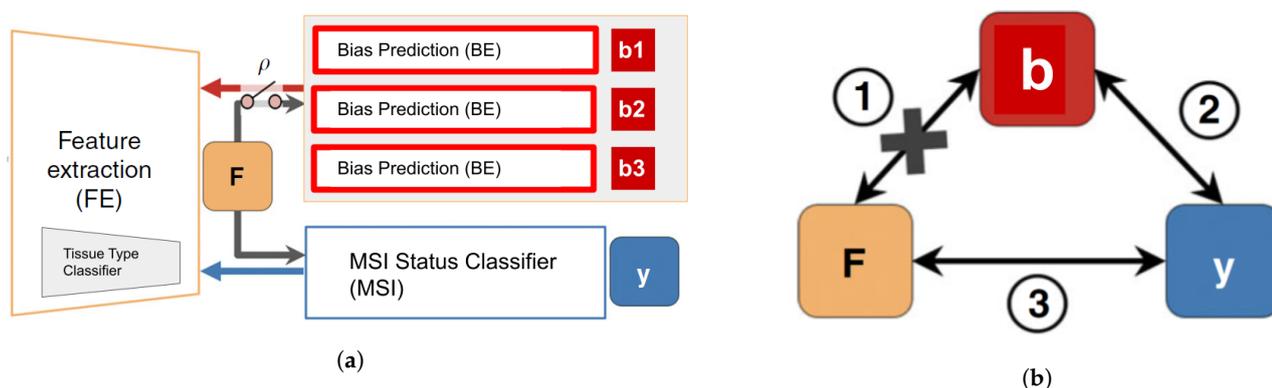


Figure 2. Network architecture: (a) The deep neural network architecture is composed of three modules, FE learns features, F , that successfully classify the input in the outcome y using MSI while being invariant (statistically independent and conditioned by ρ) to the biases variables, b_n , using the adversarial components BE and the adversarial loss. (b) The bias variables b_n , responsible for multiple batch effects, influence both the output y (i.e., ②, MSI status classification) and the input X , from which feature F is extracted (i.e., ①). The MSI classifier deems to find the relation ③ to enable prediction of the output labels while the adversarial components aim to remove the direct dependency between F and b_n . Figure adapted from [21] by renaming the modules and adding multiple adversarial components to the architecture. <http://creativecommons.org/licenses/by/4.0/>.

2.4. Bias Identification

To systematically identify and quantify the effect of potential biases interfering with the MSI prediction task, the baseline model, which as described in Section 2.3 included only the feature extractor backbone FE connected to the MSI status classifier MSI module, was trained on the study sample using 5-fold cross-validation for 3 epochs. The squared distance correlation dc [27] was computed on each batch-iteration between the extracted features from the FE and three putative biases (study project, patient, and glass) followed by backward selection to uncover hidden interactions between biases. The dc is a measure of dependence between random vectors analogous to product-moment covariance and correlation, but unlike the classical definition of correlation, distance correlation is zero only if the random vectors are statistically independent [27].

The three biases considered for testing were the study project, the patient and the TMA glass based on the following reasoning: First, since MSI-H patients were significantly over-represented in the subset of cases obtained from a single hospital (34% in HGUA vs 7.4% in EPICOLON), and also samples were obtained in different years on each project (2017 onwards in HGUA vs prior to 2013 in EPICOLON) in this study the contributing project becomes a potential task bias; i.e., prediction of MSI status may be dependent on project instead of on the image bio-markers of MSI status. Second, tiles from the same patient commonly share observable visual patterns originated from both tissue characteristics, slicing direction and staining differences in laboratory procedures. As the model is designed to be trained at tile level, even if no patient is at the same time in the training and validation set, the model would learn those patterns as shortcuts for the MSI task, hence being a potential reason for model overfitting. Overfitting was indeed verified when the baseline model was trained more than 3 epochs observing an ever decreasing

training loss near zero with no improvements in the validation loss. Third, each TMA agglutinates sample spots from tens of patients which are placed between an underlying and a cover glass. The unique characteristics associated with the digitization of each TMA-glass may be statistically associated with MSI status if the distribution of classes across TMAs is not uniform.

2.5. Bias Ablation

Once biases were identified and quantitatively characterized, the ablation of each bias b_n , was implemented through an adversary training and distillation bias regime following the approach described in [21] (see Figure 2), but differently from this work, the ablation was not limited to only one but to all the biases identified.

Namely, given the input image X , the FE module extracts a feature vector F , on top of which the MSI module predicts the class label y . To guarantee that F is not biased to the multiple b_n , each corresponding BE_n module back-propagate, in a consecutive way, the loss to FE adversarially, i.e., as $-\lambda L_{be_n}$. It results in features that minimize the classification loss of the MSI module while maintaining the least statistical dependence on each of the bias b_n .

Each of the BE modules is trained on a y -conditioned cohort, i.e., samples of the training data whose y values (MSI labels) are confined to a specific group (referred as ρ in Figure 2). Consequently, the features learned by the system are predictive of y while being conditionally independent of the batch effect originated by each of the biases. On the implemented architecture, the system would learn to separate MSI-H vs MSS samples by training each BE_n only on the MSS group to correctly model the batch effects on the samples. We perform the adversarial training of each of the BE_n only on the MSS group.

During the end-to-end system training a min-max game, is defined between two networks. The classification loss L_{msi} is defined by a cross-entropy:

$$L_{msi}(X, y; \theta_{fe}, \theta_{msi}) = - \sum_{i=1}^N \sum_{m=1}^M y_{im} \log(\hat{y}_{im}) \quad (1)$$

where X and y are the input images and corresponding msi target labels, respectively, N is the number of training pairs (X, y) , M is the number of classes to predict (two, MSI-H vs. MSS) and \hat{y} is the predicted MSI label.

Each batch effect or bias loss L_{be_n} is based on the squared Pearson correlation $corr^2$:

$$L_{be_n}(X, y; \theta_{fe}, \theta_{be_n}) = - \sum_{k=1}^K corr^2(b_k, \hat{b}_k) \quad (2)$$

where b_k defines the vector of the bias across all N training inputs. The statistical dependence is removed by pursuing a zero correlation between b_k and \hat{b}_k through adversarial training.

The overall objective of the end-to-end network is defined as:

$$\min_{\theta_{fe}, \theta_{msi}} \max_{\theta_{be}} L_{msi}(X, y; \theta_{fe}, \theta_{msi}) - \lambda \sum_{b=1}^B L_{be_n}(X, b_n; \theta_{fe}, \theta_{be_n}) \quad (3)$$

where B is the number of protected variables or biases b .

Specifically, in each iteration, first we back-propagate loss L_{msi} to update θ_{fe} and θ_{msi} . Second, for each b_n , we fix θ_{fe} and then minimize the L_{be_n} loss to update their corresponding θ_{be_n} . Finally, we fix θ_{be_n} and then maximize the L_{be_n} loss to update θ_{fe} , hence distilling all the biases from θ_{fe} . In this study, each L_{be_n} depends on the correlation operation, which is a population-based operation, as opposed to individual level error metrics such as cross-entropy or MSE losses. Therefore, we calculate the correlations over each training batch as a batch-level operation. In conclusion, FE extracts features that minimize the classification criterion, while 'fooling' all BE modules (i.e., making each BE incapable of

predicting their corresponding bias). Hence, the saddle point for this objective is obtained when the parameters θ_{FE} minimize the classification loss L_{msi} while maximizing the L_{ben} loss of all the BE_n modules.

2.6. Partitions, Model Training and Metrics

For model training, the dataset was split 80/20 for training and validation applying 5-fold cross-validation and guaranteeing on each fold that the images of each patient only belonged to one set (either training or validation but not both). To address imbalance in MSI status as well as in the number of tiles for each patient, we applied a composite weighted random sampling for both criteria, resulting in a balanced training set for both patient and MSI label simultaneously. Tiles were resized to 224×224 pixels and color was normalized following Mazenko method [24]. In addition to Mazenko, experiments were done with and without an additional color normalization with statistics computed from the EPICOLON image dataset which included 25 different hospitals. Data augmentation at training time consisted in random rotations up to 90° , dihedral flips with probability of 0.5, a perspective warping of maximum 0.2, and hue variations of maximum 0.15. Of note, training sets included all magnifications so that the network could be trained simultaneously on higher tissue architecture patterns as well as on cellular-level features including nuclear characteristics. The batch size was 512 images.

The statistical dependency between learned features and each of the selected biases was monitored during model training with the squared distance correlation dc . Principal component analysis (PCA) was used to assess how the spatial representations of the learned features were affected by the protected (bias) variables before and after bias ablation. One-way ANOVA was used to compare the false positive rates and false negative rates of the different tissue types and magnifications on the MSI classification task. Metrics used to assess the performance of the MSI classifier include AUC, balanced accuracy, sensitivity, specificity, positive predictive value, negative predictive value, false positive rate and false negative rates. Metrics dependent on MSI-H prevalence are calculated assuming 15% in the real population.

2.7. Explainable Methods

SHAP (SHapley Additive exPlanations) [28] values were used to provide a means of visually interpreting the topology and morphology of features that influence predictions. The goal of SHAP is to explain the prediction of an instance by computing the contribution of each feature to the prediction. The SHAP explanation method computes Shapley values from coalitional game theory.

3. Results

3.1. Experiment 1: Tissue Classifier

The tissue classifier module reached an AUC of 0.98 in the validation set. This module was capable of classifying the different regions of the image (tumor epithelium, stroma, normal epithelium, mucin, muscular fibers, lymphocytic infiltrates, debris, adipose tissue and background) as shown in Figure 3.

After the automatic filtering done by the preprocessing module, the final study sample totaled 1788 patients (171 MSI-High). The image dataset consisted of 1,065,479 tiles or adjacent regions without overlap including all magnifications and all types of tissues. The tissue classifier module was then used in inference for the pre-selection of the regions of interest restricted to tumor epithelium, lymphocytic infiltrates and mucin totaling 523,624 tiles, linked to the entrance to the MSI module classifier. On this filtered dataset, the distribution of tiles by each magnification is shown in Table 2.

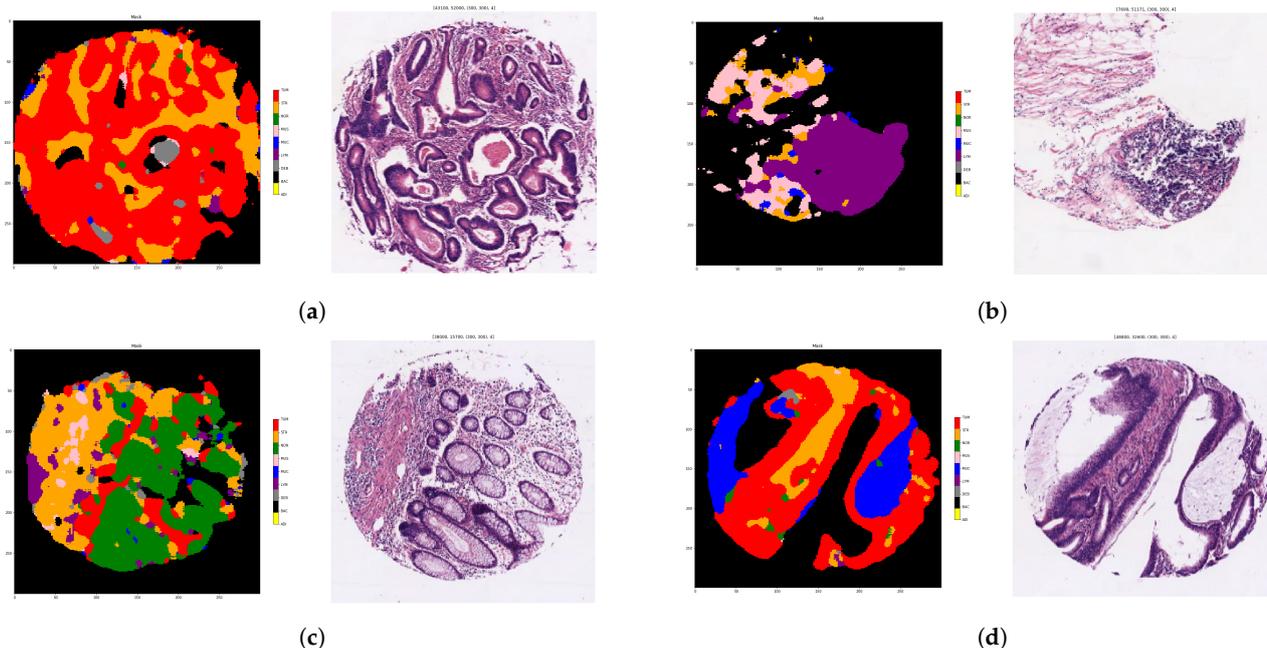


Figure 3. Examples of spots with each tissue region coded with different colors: tumor epithelium (red), stroma (orange), normal epithelium (green), mucin (blue), muscular fibers (pink), lymphocytic infiltrates (purple), debris (grey), adipose tissue (yellow) and background (black). Each sub-figure (a–d), corresponds to one example of spot where the original sample is placed on the right and the corresponding color mask is placed on the left. The tissue class for each pixel is computed as the majority tissue class of partial overlapping tiles applying a sliding window to raster each spot. Tiles were input to the tissue classifier module at $\times 20$ magnification.

Table 2. Distribution of tiles by magnification. Tiles were filtered by regions of interest restricted to tumor epithelium, lymphocytic infiltrates and mucin.

Magnification	Number of Tiles
$\times 40$	386,524
$\times 20$	100,569
$\times 10$	23,630
$\times 5$	9954
$\times 0$	2947
All (total)	523,624

3.2. Experiment 2: Bias Identification, Interaction between Variables and Bias Ablation Implementation

First a baseline model was trained for 10 epochs and overfitting was observed with an ever decreasing training loss approaching zero with no improvements in the validation loss. The lack of generalizability in validation was suspected to the presence of biases described in Section 2.

Second, a baseline model was trained for 3 epochs, with early stopping, and the statistical dependence between the representations learned by the feature extractor backbone with regard to the target task (MSI classification) and each of the suspected bias was measured in the entire training cohort with the squared distance correlation dc . As shown in Table 3, the learned features by the baseline network had surprisingly a higher dc with the project variable (0.17) when compared with the MSI label (0.08). Subsequently, as a subgroup analysis, dc was computed for each project cohort separately. At this step a hidden interaction emerged between the HGUA project cohort and the TMA glass and patient variables respectively, where the dc increased up to 0.29 in both cases. The unexpectedly strong association between the TMA glass and the learned representations only occurring in the HGUA project was further investigated and it obeyed to the MSI label distribution of the spots among TMAs. Specifically, we found that many TMAs in the HGUA project

had spots from only one class (either MSS or MSI-H, but not both). Moreover, as the glass covered the spots it meant that not only the background but also the tissue regions carried the glass characteristic patterns. This observation was relevant implying that removing background regions -even if necessary- is not enough to remove this bias. In contrast, all of the EPICOLON TMAs had a varying representation of spots of both classes, though always strongly unbalanced towards the MSS class.

Third, once confirmed the presence of multiple bias variables (i.e., project, patient, and TMA glass variables), three corresponding bias-rejecting learners were attached to the end-to-end network applying an adversarial consecutive training regime for each head as explained in Section 2.3. After implementing the bias-ablation technique and retraining, we recomputed the statistical dependence between the representations learned by the bias-ablated model with respect to the target and each bias. As shown in Table 3, associations of the learned features with the MSI status were strengthened, increasing from 0.08 to 0.21 in the overall study set. Conversely, the dependence with three biases was weakened, decreasing from 0.17 to 0.02 with respect to the project bias and from 0.08 to 0.02 in the case of the patient and glass bias. When narrowing down to the HGUA project, the bias ablation resulted in a higher association of the learned features with the MSI status, but failed to disentangle their association with the glass and patient bias which still remained high. As explained above, in the HGUA project some TMAs were composed only of one single target class of spots, hence, the observed persistent statistical dependence supports that the attempt to decoupling those TMA glasses from the MSI status is unfeasible and that the glass may provide the network with an unfair shortcut for MSI status prediction.

Table 3. Statistical dependence between the learned features with regard to the target task (MSI) and each bias quantitatively measured by the squared distance correlation (dc). Hidden interactions between the different biases are explored by subgroups by re-calculating the dc for each project. The learned features from a baseline model are compared against a bias-ablated model with regard to their statistical dependency to the target task and each bias. In bold are highlighted the lowest distance correlations achieved for each bias. The bias-ablated model maximizes the association of the learned features with respect to the MSI task and minimizes the statistical mean dependence with the biases.

Model	Subgroup	MSI dc	Project Bias dc	Patient Bias dc	Glass Bias dc
Baseline	All	0.08	0.17	0.08	0.07
Baseline	EPICOLON Project	0.04	-	0.02	0.07
Baseline	HGUA Project	0.46	-	0.29	0.29
Bias-ablated	All	0.21	0.02	0.02	0.02
Bias-ablated	EPICOLON Project	0.15	-	0.01	0.07
Bias-ablated	HGUA Project	0.62	-	0.35	0.33

3.3. Experiment 3: Bias Ablation Analysis

Statistical dependence, conditioned in the MSS cohort as explained in Section 2, between the learned features and each bias measured with the squared distance correlation was consistently reduced by more than half in the bias-ablated models as compared to the baseline models, as shown in Table 4. The dc with the project bias was the most reduced one. During model training, as shown in Figure 4, the squared distance correlation between the learned features and the target MSI status (blue) increased as expected as a function of the number of training iterations. When comparing the two training regimes, we observed that in the baseline model (Figure 4a) the correlation with the MSI target increased together with all other identified biases and especially with the sample project (orange) which overlapped with the MSI target (blue). In contrast, in the model that applied the bias-ablation adversarial training (Figure 4b), the dc between the learned features and the project and patient biases did not increase, maintaining them at minimum values, while

simultaneously allowing the increase of statistical dependence with the MSI target. Notably, the efficacy of the bias ablation was more marked with regard to the project bias (orange).

Table 4. Statistical dependence between the learned features and each bias for the MSS group. Quantitatively measured by the squared distance correlation (dc). The mean and standard deviation of the dc obtained from the 5 fold baseline models is compared against the 5 fold bias-ablated models. In bold are highlighted the lowest distance correlations achieved for each bias.

Model	Project Bias	Patient Bias	Glass Bias
Baseline model	0.25 ± 0.05	0.12 ± 0.02	0.07 ± 0.008
Bias-ablated model	0.10 ± 0.04	0.04 ± 0.018	0.03 ± 0.013

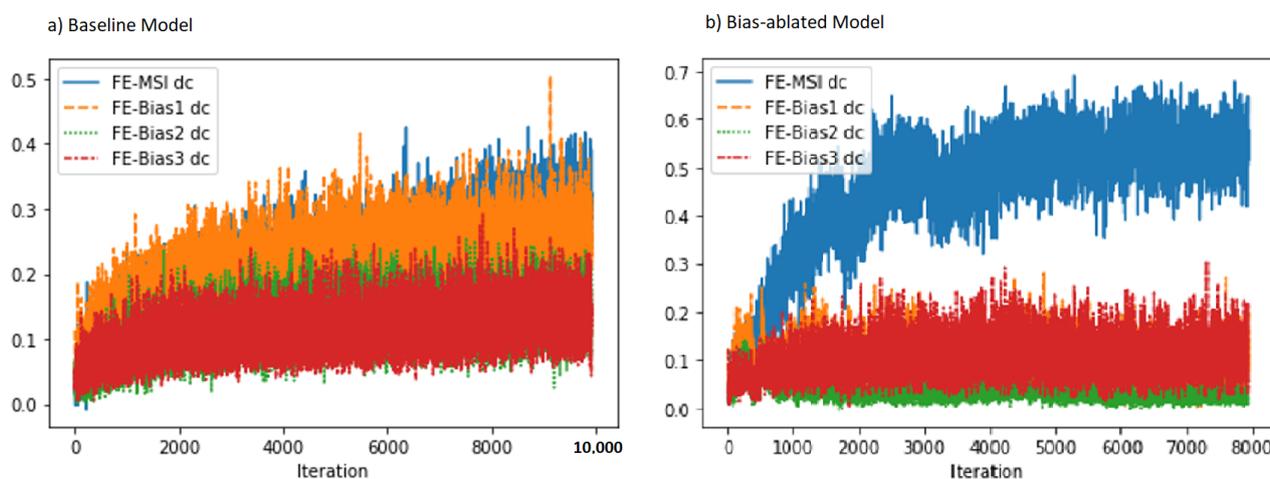


Figure 4. Statistical dependency, measured with the squared distance correlation (dc), between the learned features, MSI learning task and each bias (project, patient and glass bias) during the training process. After bias-ablation, the statistical dependency of the learned features with regard to MSI-classification is increased (blue), while reduced in particular with regard to the project bias (orange). Fe-Bias1 dc = Squared distance correlation between learned features and project bias (orange). Fe-MSI dc = Squared distance correlation between learned features and MSI status (blue). Fe-Bias2 dc = Squared distance correlation between learned features and patient bias (green). Fe-Bias3 dc = Squared distance correlation between learned features and TMA glass bias (red).

The PCA projection of the learned representations comparing the baseline model vs the bias-ablated model is shown in Figure 5. The bias ablation technique results in a better representation space which is more invariant to the bias variables (project and TMA glass) while the baseline shows patterns more influenced by the bias. For example, and as illustrated in Figure 5c, most of the TMA glasses are organized as spatial clusters in the baseline model which means that the representation space of the learned features is not invariant with regard to the TMA glass. Conversely, after applying the bias-ablation technique colors of TMA glasses are less organized and follows a more homogeneous distribution Figure 5d.

To assess the impact of the bias-ablation technique on the MSI classification performance, we compared the validation results of 5-folds models trained for up to 3 epochs with and without the adversary ablation of the three known biases. As shown in Figure 6, when models were evaluated on the validation set without differentiating between the projects of samples (i.e., both EPICOLON and HGUA samples included), there was not appreciable impact on performance between the bias-ablated vs baseline models at image level (both achieved an AUC of 0.87 ± 0.03). Nonetheless, when the baseline models were evaluated separately for each of the different projects (EPICOLON vs HGUA), the performance in HGUA samples (AUC = 0.97 ± 0.01 , balanced accuracy = 0.91 ± 0.03) was significantly higher than in EPICOLON samples (AUC = 0.82 ± 0.03 , balanced accu-

racy = 0.75 ± 0.03) with a gap of up to 0.15 and 0.16 in the AUC and balanced accuracy respectively. This project gap was attenuated to 0.13 (in both the AUC and balanced accuracy) after applying the bias-ablation technique, by decreasing the performance in HGUA (AUC = 0.96 ± 0.02 , balanced accuracy = 0.89 ± 0.02) while increasing the performance in EPICOLON (AUC = 0.83 ± 0.03 , balanced accuracy = 0.76 ± 0.01).

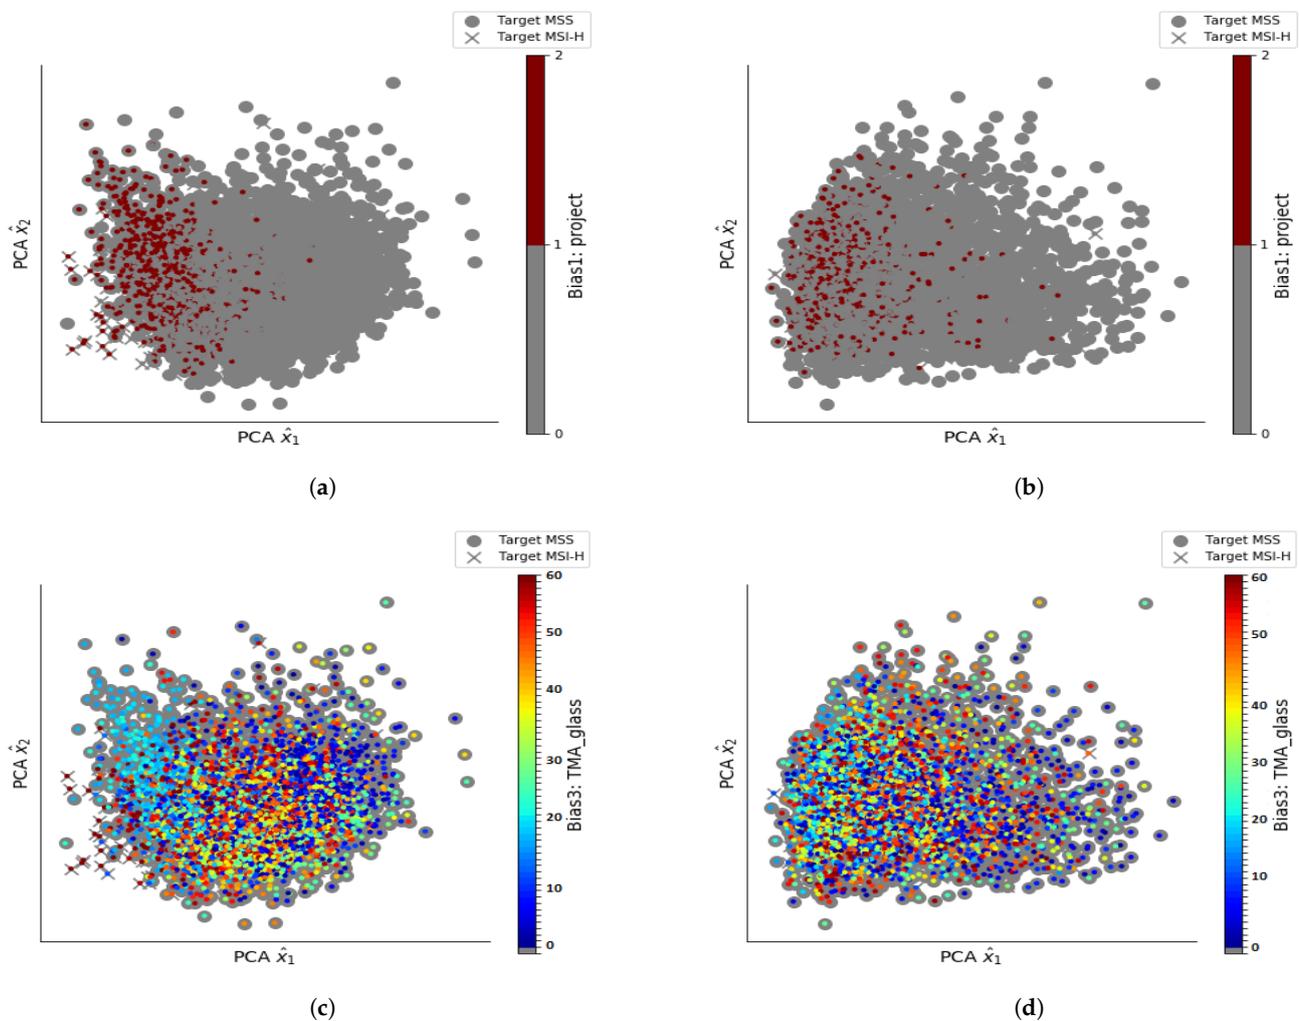


Figure 5. PCA projection of the learned features for baseline vs bias-ablated models. Color encodes the category for project bias (a,b) and TMA glass bias (c,d). (a) Baseline model features and project bias. (b) Bias-ablated model features and project bias. (c) Baseline model features and TMA glass bias. (d) Bias-ablated model features and TMA glass bias.

3.4. Experiment 4: Results at Tile vs. Patient Level and Effect of Tissue Types and Magnifications

For each of the 5 folds, tile predictions of the bias-controlled model were aggregated by majority voting for the decision of the MSI status at the patient level. 5-folds validation sets consisted of approximately 300 patients each and the mean prevalence of MSI-H was 9%. In order to simulate performance on real population prevalence, metrics dependent on disease prevalence (see Table 5) were calculated assuming a MSI-H prevalence of 15%.

The AUC at tile level, and including all three selected tissues (tumor epithelium, mucin and lymphocytic regions) and all 4 magnifications, was 0.87 ± 0.03 and increased to 0.9 ± 0.03 at patient level.

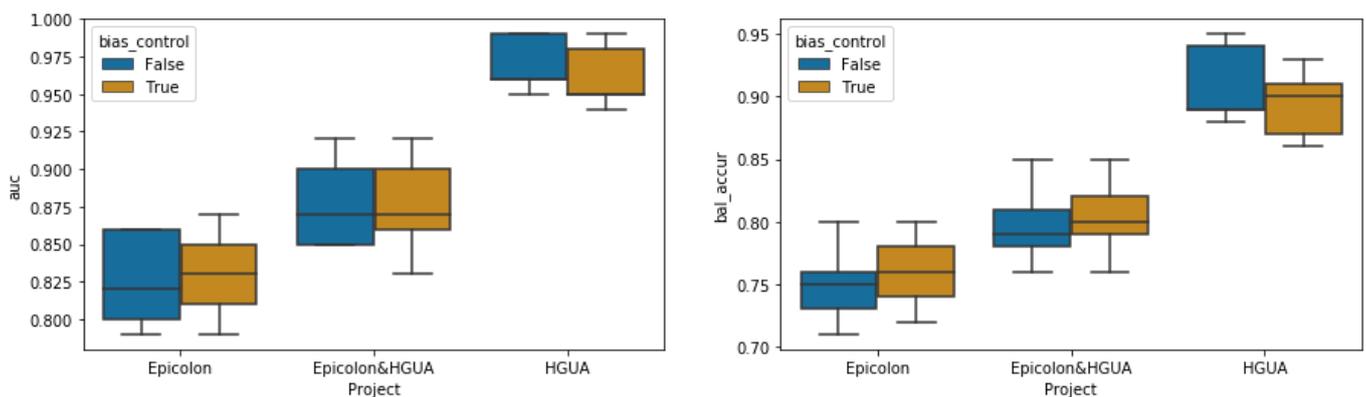


Figure 6. MSI classification results at tile level comparing the 5-fold bias-ablated models (orange) vs. 5-fold baseline models (blue) across the different projects (EPICOLON and HGUA).

The AUC at tile level for a bias-controlled model trained only on tumor epithelium was 0.87, but differently from the model that included all three selected tissues the AUC at patient level did not increase and remained 0.87. In the models trained with all three tissue types, tiles tagged as lymphocytic infiltrates (LYM) and mucin (MUC) had an average of false positive rate (ratio of false positive over true negatives) of 0.32 and 0.22 respectively, compared to 0.12 in tumoral epithelium (TUM) tiles (see Figure 7). This finding supports that nuclear and cellular characteristics of tumor epithelium are more specific of MSI-H status, while mucin and lymphocytic infiltrates are nonspecific, which would explain the higher rates of false positives at tile level. Nonetheless, as inferred from the performance gain in AUC at patient level, regions with mucin and lymphocyte's infiltrates were probed to be relevant for MSI classification at case level by contributing to a lower rate of false negatives as shown in Figure 7. Because MSI-H tumors typically exhibit larger amounts of mucin and lymphocytic infiltrates, even if not specific, it would explain the increase in model performance at patient level after applying the majority of votes. We compared the effect of the different types of tissues (TUM, LYM, MUC) on the false positive rates and false negative rates. A one-way ANOVA found a statistically significant difference in the false positive rates between at least two groups ($F(2, 12) = 165.2, p < 0.001$). Tukey's HSD Test for multiple comparisons found that the mean value of the false positive rate was significantly different between all combinations of groups ($p < 0.01$). Similarly a one-way ANOVA found a statistically significant difference in the false negative rates between at least two groups ($F(2, 12) = 58.9, p < 0.001$). Tukey's HSD Test for multiple comparisons found that the mean value of the false negative rate was significantly different only between TUM and LYM ($p < 0.01$) and between LYM and MUC ($p < 0.01$).

The AUC at tile level for bias-controlled models trained respectively at $\times 5$, $\times 10$, $\times 20$ and $\times 40$ magnifications in all three selected tissue types were 0.81, 0.85, 0.87, 0.86. According to the AUC results and as shown in Figure 7, the specificity was higher for $\times 20$ magnification where the lowest rates of false positives were observed. Nonetheless false negative rates decreased (lower cases missed) with decreasing magnification from $\times 40$, $\times 20$ up to $\times 10$ magnification, which argues in favor that including high level tissue architectural patterns (up to $\times 10$) contribute to increased sensitivity. The AUC at tile level for bias-controlled models and including all available magnifications was 0.87, hence the inclusion of all magnifications was non-inferior to including only one single magnification for training, and was also considered as a valuable data augmentation technique.

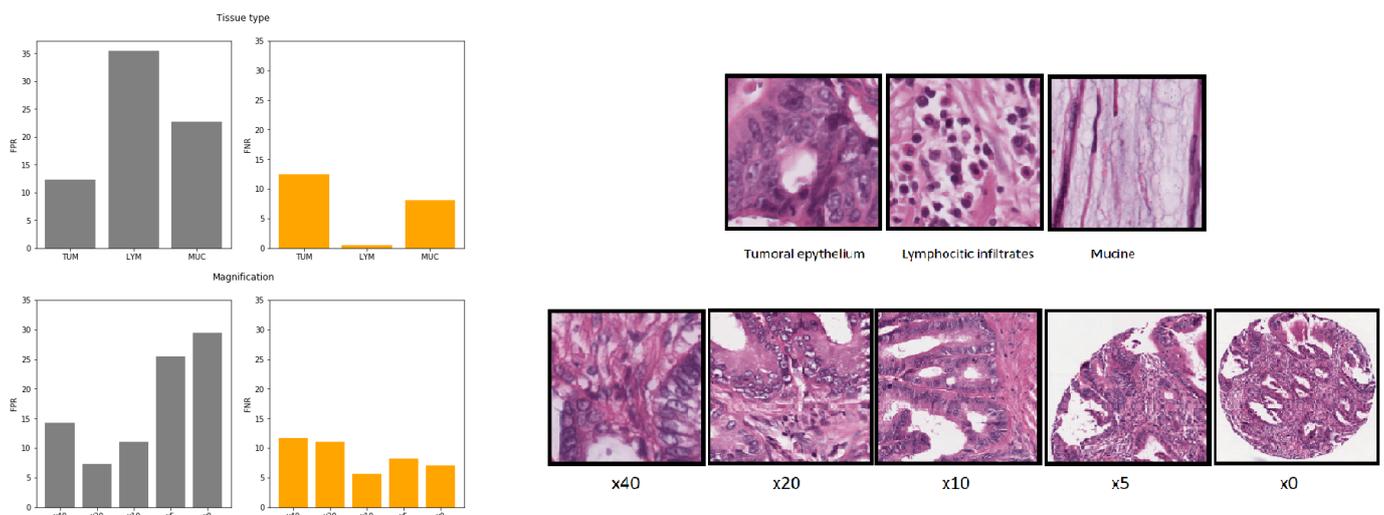


Figure 7. Average of false positive rates (gray) and false negative rates (orange) at tile level for the 5 fold bias-controlled models by tissue type (figures above): tumor epithelium (TUM), lymphocytic infiltrates (LYM) and mucin (MUC) and different magnifications (figures below): $\times 40$, $\times 20$, $\times 10$, $\times 5$, $\times 0$.

Table 5. Patient level performance of the bias-ablated models cross-validated in 5 folds. Metrics that are dependent on MSI-H prevalence, are calculated assuming a MSI-H prevalence in the real population of 15%. Confidence intervals for sensitivity, specificity and accuracy are “exact” Clopper-Pearson confidence intervals. Confidence intervals for the predictive values are the standard logit confidence intervals given by [29]. P = Prevalence, S = Sensitivity, E = Specificity. * Metrics that are dependent on MSI-H prevalence.

Metric	Value	95% CI	Definition
AUC	0.90	0.87–0.93	
Accuracy *	88%	83.8–91.4%	$S * P + E * (1 - P)$
Sensitivity (S)	87%	79.5–91.3%	$\frac{TP}{TP+FN}$
Specificity (E)	88.3%	86.5–89.9%	$\frac{TN}{TN+FP}$
Positive Predictive Value *	56.5%	52.6–60.4%	$\frac{S * P}{S * P + (1 - E) * (1 - P)}$
Negative Predictive Value *	97.3%	96–98.2%	$\frac{E * (1 - P)}{E * (1 - P) + (1 - S) * P}$

We compared the effect of the different magnifications ($\times 40$, $\times 20$, $\times 10$, $\times 5$, $\times 0$) on the false positive rates and false negative rates. A one-way ANOVA found a statistically significant difference in the false positive rates between at least two groups ($F(4, 20) = 53.6$, $p < 0.001$). Tukey’s HSD Test for multiple comparisons found that the mean value of the false positive rate was significantly different between all combinations of groups ($p < 0.01$) except between the pairs of magnification $\times 40$ – $\times 10$, $\times 20$ – $\times 10$ and $\times 5$ – $\times 0$. Similarly a one-way ANOVA found a statistically significant difference in the false negative rates between at least two groups ($F(4, 20) = 8.2$, $p < 0.001$). Tukey’s HSD Test for multiple comparisons found that the mean value of the false negative rate was significantly different ($p < 0.01$) only between the pairs of magnification $\times 40$ – $\times 10$, $\times 40$ – $\times 0$ and $\times 20$ – $\times 10$.

3.5. Experiment 5: Visual Explainability

For visually interpreting the topology and morphology of features that influenced the predictions, Figure 8 shows that the most relevant features for classifying the samples as MSI-high vs MSS were located at conglomerates of tumoral cells with activations mainly gathered at the nucleus of the cells. Note that as a binary classification task, for the MSI-H projection images the red pixels are the most important regions influencing the prediction

towards the MSI-H class, while for the MSS image projection the same pixels are shown as blue, being symmetrical.

When comparing the projections of SHAP values between the baseline and the bias-ablated models for the same inputs, there were no perceptible differences by visual observation, nonetheless the discriminative power measured by the softmax values distance between classes is increased, in particular for MSS samples.

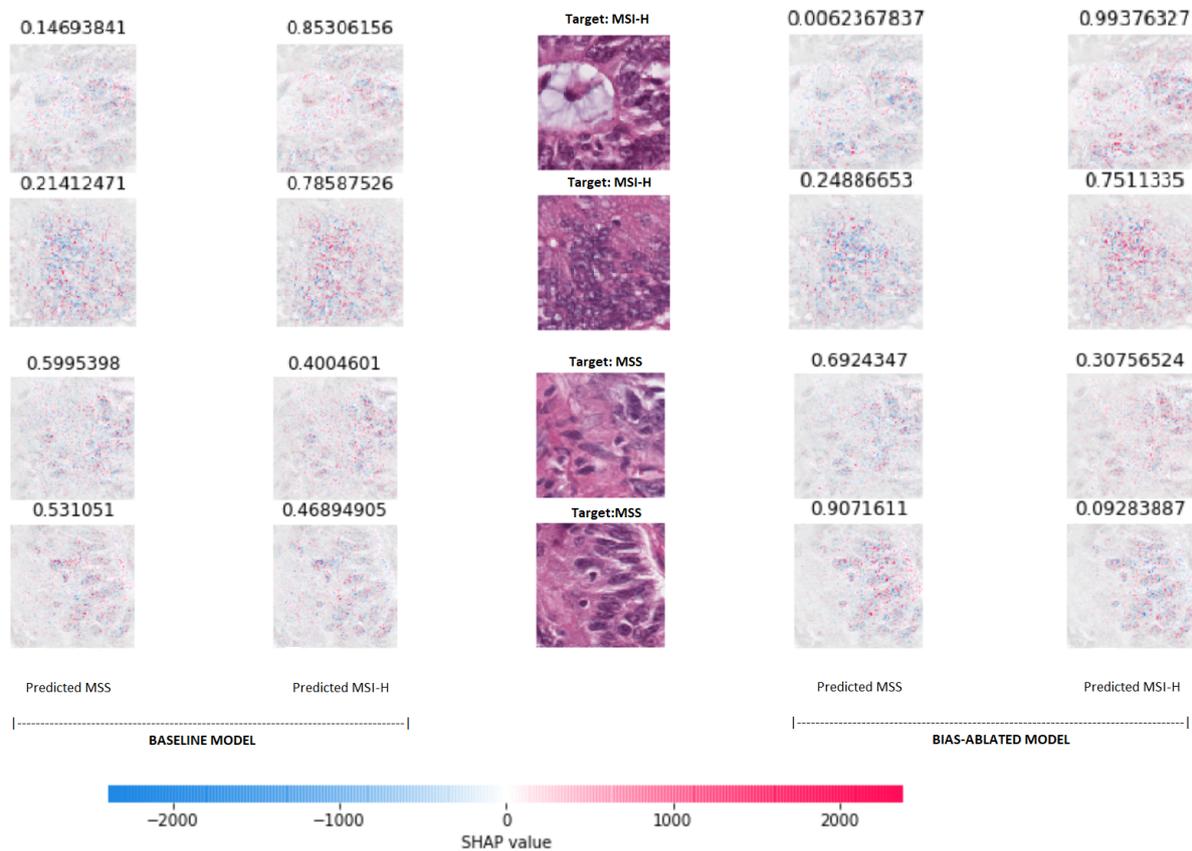


Figure 8. SHAP (SHapley Additive exPlanations) values projected on the tiles to visually interpreting the topology and morphology of features that influence predictions of the baseline and bias-ablated model (left and right respectively). The input samples are placed in the center of the graph with their respective target labels (2 MSI-H samples at the top and 2 MSS samples below). For each of the input tiles, the resulting SHAP values obtained for both models are projected for each class on top of a grayscale copy of the input tile (left and right for MSS and MSI-H respectively). Red pixels increase the model's output while blue pixels decrease the output. The sum of the SHAP values equals the difference between the expected model output (averaged over the background dataset) and the current model output. The softmax values (prediction) for each class (MSS vs. MSI-H) are shown on top of each projection.

4. Discussion

We present a system for the prediction of MSI from H&E images using artificial vision techniques that incorporates an end-to-end TMA-customized image preprocessing module to tile samples at multiple magnifications in the regions of interests guided by the automatically detected type of tissues and a multiple bias distiller system integrated with the MSI predictor.

In the present work we find that TMAs have special characteristics, not reported for WSIs, which make them especially challenging for the application of DL methods in digital pathology, emphasizing the relevance of addressing biases.

A systematic study of biases at tile level demonstrated three hidden variables interfering with the model's learned representations: the project of origin of samples, the patient's spot and the TMA glass where each spot was placed. Even if it is preferred to control

for any of those types of biases at the dataset level, it entails either obtaining more tissue and/or re-allocating them in new TMA glasses which was unfeasible in general. Instead, we reused the TMAs as they were provided for research purposes given that first, the most optimal management of tissue samples avoiding sample waste is always desirable and second, the presence of associations spurious or otherwise undesirable that are exploited by DL models, rather than being an odd problem affecting only our work, is a general, common and not yet resolved challenge in medical datasets used to train AI systems. Consequently, we decided to dedicate the efforts to systematically study and address the biases at the learning stage. For this purpose, a novel multiple bias rejecting technique has been implemented at the deep learning architecture to directly avoid learning the batch effects introduced by protected variables.

The implemented method achieved a significant reduction in the dependence of the learned features with regard to the project bias and patient bias in the general study populations but did not reduce the glass bias dependence in the HGUA project where we found that for most of the TMAs there were only samples with one single type of target class included. We observed that the ablation method is highly effective for mitigating bias in datasets where for all possible ordered pairs of protected variable classes and target classes there are representative samples even if heavily in-balanced. In addition, the performance in the MSI classification improved in terms of AUC and balanced accuracy for the population meeting this condition (EPICOLON cohorts). Conversely, when this condition was not met (as in the case of the HGUA for the glass bias, where the target and protected variable had an unequivocal association in the samples), the statistical dependence was not eliminated and the classification performance in the target task still seemed to exploit the bias maintaining a highly marked predictive advantage in comparison with the classification results in the EPICOLON cohorts. We observed that the performance of the models were tightly associated to the presence of batch effects and their presence had a larger impact on performance than the number of patients available for training on each cohort. In essence, as illustrated in Figure 6, even if the number of patients was smaller in the HGUA cohort, the model performance on this cohort was higher than that achieved in the EPICOLON project, which was explained by the model's exploitation of the TMA-glass bias in the HGUA cohort. When analyzed considering all study population, the learned features from the bias-ablated model had maximum discriminative power with respect to the task and minimal statistical mean dependence with the biases.

In contrast to other population-based cohorts where MSI prevalence is around 15%, in the EPICOLON project, the cohorts had a mismatch repair deficiency in only 7.4% patients. In EPICOLON I only 91 patients (7.4%) had a mismatch repair deficiency with tumors exhibiting either genetic microsatellite instability ($n = 83$) or loss of protein expression ($n = 81$) [2]. This difference is attributed to the fact that the EPICOLON project was a population-based study; while CRC cohorts with a 14–15% dMMR prevalence usually correspond to registries with a higher percentage of patients with family history, as it is for example, the Cancer Family Registry [1].

Regarding the classification MSI status at case level, we observed that the performance consistently increased in all experiments when not only tumor epithelium but also the mucinous and lymphocytic infiltrate regions were included. Those regions were probed to be nonspecific at image level, but increased the sensibility at patient level, which is a desirable characteristic for screening purposes.

Also, a $\times 20$ magnification achieved the higher specificity, but at the same time, reducing magnifications up to $\times 10$ contributed to higher sensitivity of the models. This observation supports that the best approach would be to include different magnifications, helping the model to learn both low and high level tissue architectural patterns at the same time. Moreover, including all magnifications during training was considered in this project as a data augmentation technique.

The tissue classifier module reached an AUC of 0.98 in the validation set. This module was capable of classifying the different regions of the image: tumor epithelium, stroma,

normal epithelium, mucin, muscular fibers, lymphocytic infiltrates, debris, adipose tissue and background. Regarding the final task of MSI status prediction, the AUC at tile level, including all three selected tissues (tumor epithelium, mucin and lymphocytic regions) and all magnifications, was 0.87 ± 0.03 and increased to 0.9 ± 0.03 at patient level.

Limitations of the study are as follow: We observed, after applying the image pre-processing pipeline, a reduction in the final available MSI population, altering its original frequency in the EPICOLON project. Specifically, due to lack of epithelial tumor regions in the spots, up to 9% ($n = 159$ patients) in the EPICOLON population were excluded from the final analysis, where the largest proportion of excluded patients corresponded to the MSI-H arm, hence reducing its frequency from the expected 7.4% in EPICOLON to 5.8% ($n = 105$). This was explained by a more intensive molecular tissue testing performed in MSI-H cases in the context of the EPICOLON research project that exhausted a subset of spots with no viable regions of epithelial tumor left. The reduction in the final available MSI population, altering its original frequency, was consequently not longer considered a representative sample of the original EPICOLON MSI-H population. To overcome this limitation, the dataset was further enriched with cases from a single hospital as shown in Figure 1. On the one hand, this enrichment would increase the exposure of the MSI classifier to additional unselected MSI-H cases and, on the other hand, the bias rejecting technique implemented successfully addressed the batch effect introduced by the project of origin. Finally, result metrics which are impacted by disease prevalence, are calculated considering the MSI-H prevalence in the real population (15%), so as to approximate its performance in the clinical setting.

As future work, only after addressing the remaining TMA-glass bias at the data level, testing the generalizability of the system in an independent and prospective test would be necessary. Also multimodal variables including age, stage, location, Bethesda criteria could be included in the model to explore their potential to improve the predictive capacity.

5. Conclusions

In this study, we present an AI system for the prediction of MSI in colorectal cancer from H&E images in TMAs. The AUC at tile level, and including all three selected tissues (tumor epithelium, mucin and lymphocytic regions) and 4 magnifications, was 0.87 ± 0.03 and increased to 0.9 ± 0.03 at patient level. The system incorporates a tissue type classifier module to select the regions of interest and a multiple bias rejecting technique based on adversarial training. The learned features from the bias-ablated model have maximum discriminative power with respect to the task of MSI prediction and minimal statistical mean dependence with the biases. We emphasize the need of analyzing biases exploited by DL system in digital pathology and propose a method to address multiple biases at the DL architecture.

Author Contributions: Conceptualization, C.A. and A.B. and A.P.; methodology, A.B. and C.A. and A.T.; software, A.B. and A.T.; validation, A.B. and C.A.; formal analysis, A.B.; investigation, A.B. and A.P. and A.T. and R.J. and X.L. and X.B. and A.C. and Á.C. and C.A.; resources, C.A. and A.B.; data curation, C.A. and R.J. and X.L. and X.B. and A.C. and Á.C.; writing—original draft preparation, A.B. and C.A.; writing—review and editing, A.B. and A.P. and A.T. and R.J. and X.L. and X.B. and A.C. and Á.C. and C.A.; visualization, A.B.; supervision, C.A.; project administration, C.A.; funding acquisition, C.A. All authors have read and agreed to the published version of the manuscript.

Funding: This work and APC was supported by the Programa Estatal de Generación de Conocimiento y Fortalecimiento del Sistema Español de I+D+i, funded by the Instituto de Salud Carlos III and Fondo Europeo de Desarrollo Regional 2014-2020 (Grant DTS19/00178).

Institutional Review Board Statement: The study was conducted according to the guidelines of the Declaration of Helsinki, and approved by the Institutional Review Board of ISABIAL (CEIm PI2019-029).

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: Image tiles for the NCT-CRC-HE-100K and CRC-VAL-HE-7K datasets are available online <https://zenodo.org/record/1214456#.XcNpCpNKjyw> (accessed on 25 October 2019).

Acknowledgments: We want to particularly thank the patients and the BioBank ISABIAL integrated in the Spanish National Biobanks Network and in the Valencian Biobanking Network for their collaboration.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

<i>BE</i>	batch effect module learners
<i>FE</i>	feature extractor backbone
<i>MSI</i>	Microsatellite Instability module learner
<i>dc</i>	squared distance correlation

References

1. Moreira, L.; Balaguer, F.; Lindor, N.; De la Chapelle, A.; Hampel, H.; Aaltonen, L.A.; Hopper, J.L.; Le Marchand, L.; Gallinger, S.; Newcomb, P.A.; et al. Identification of Lynch syndrome among patients with colorectal cancer. *JAMA* **2012**, *308*, 1555–1565. [[CrossRef](#)]
2. Piñol, V.; Castells, A.; Andreu, M.; Castellví-Bel, S.; Alenda, C.; Llor, X.; Xicola, R.M.; Rodríguez-Moranta, F.; Payá, A.; Jover, R.; et al. Accuracy of revised Bethesda guidelines, microsatellite instability, and immunohistochemistry for the identification of patients with hereditary nonpolyposis colorectal cancer. *JAMA* **2005**, *293*, 1986–1994. [[CrossRef](#)]
3. West, N.P.; Gallop, N.; Kaye, D.; Glover, A.; Young, C.; Hutchins, G.G.; Brockmoeller, S.F.; Westwood, A.C.; Rossington, H.; Quirke, P.; et al. Lynch syndrome screening in colorectal cancer: Results of a prospective two-year regional programme validating the NICE diagnostics guidance pathway across a 5.2 million population. *Histopathology* **2021**. [[CrossRef](#)]
4. Nunes, L.; Aasebø, K.; Mathot, L.; Ljungström, V.; Edqvist, P.H.; Sundström, M.; Dragomir, A.; Pfeiffer, P.; Ameer, A.; Ponten, F.; et al. Molecular characterization of a large unselected cohort of metastatic colorectal cancers in relation to primary tumor location, rare metastatic sites and prognosis. *Acta Oncol.* **2020**, *59*, 417–426. [[CrossRef](#)] [[PubMed](#)]
5. Jenkins, M.A.; Hayashi, S.; O’Shea, A.M.; Burgart, L.J.; Smyrk, T.C.; Shimizu, D.; Waring, P.M.; Ruzskiewicz, A.R.; Pollett, A.F.; Redston, M.; et al. Pathology features in Bethesda guidelines predict colorectal cancer microsatellite instability: A population-based study. *Gastroenterology* **2007**, *133*, 48–56. [[CrossRef](#)] [[PubMed](#)]
6. Kakar, S.; Burgart, L.J.; Thibodeau, S.N.; Rabe, K.G.; Petersen, G.M.; Goldberg, R.M.; Lindor, N.M. Frequency of loss of hMLH1 expression in colorectal carcinoma increases with advancing age. *Cancer Interdiscip. Int. J. Am. Cancer Soc.* **2003**, *97*, 1421–1427.
7. Sepulveda, A.R.; Hamilton, S.R.; Allegra, C.J.; Grody, W.; Cushman-Vokoun, A.M.; Funkhouser, W.K.; Kopetz, S.E.; Lieu, C.; Lindor, N.M.; Minsky, B.D.; et al. Molecular biomarkers for the evaluation of colorectal cancer: Guideline from the American Society for Clinical Pathology, College of American Pathologists, Association for Molecular Pathology, and American Society of Clinical Oncology. *Am. J. Clin. Pathol.* **2017**, *147*, 221–260. [[CrossRef](#)]
8. Shaikh, T.; Handorf, E.A.; Meyer, J.E.; Hall, M.J.; Esnaola, N.F. Mismatch repair deficiency testing in patients with colorectal cancer and nonadherence to testing guidelines in young adults. *JAMA Oncol.* **2018**, *4*, e173580. [[CrossRef](#)]
9. Hildebrand, L.A.; Pierce, C.J.; Dennis, M.; Paracha, M.; Maoz, A. Artificial Intelligence for Histology-Based Detection of Microsatellite Instability and Prediction of Response to Immunotherapy in Colorectal Cancer. *Cancers* **2021**, *13*, 391. [[CrossRef](#)]
10. Zhang, R.; Osinski, B.L.; Taxter, T.J.; Perera, J.; Lau, D.J.; Khan, A.A. Adversarial deep learning for microsatellite instability prediction from histopathology slides. In Proceedings of the 1st Conference on Medical Imaging with Deep Learning (MIDL 2018), Amsterdam, The Netherlands, 4–6 July 2018; pp. 4–6.
11. Kather, J.N.; Pearson, A.T.; Halama, N.; Jäger, D.; Krause, J.; Loosen, S.H.; Marx, A.; Boor, P.; Tacke, F.; Neumann, U.P.; et al. Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. *Nat. Med.* **2019**, *25*, 1054–1056. [[CrossRef](#)] [[PubMed](#)]
12. Kather, J.N.; Heij, L.R.; Grabsch, H.I.; Loeffler, C.; Echle, A.; Muti, H.S.; Krause, J.; Niehues, J.M.; Sommer, K.A.; Bankhead, P.; et al. Pan-cancer image-based detection of clinically actionable genetic alterations. *Nat. Cancer* **2020**, *1*, 789–799. [[CrossRef](#)] [[PubMed](#)]
13. Cao, R.; Yang, F.; Ma, S.C.; Liu, L.; Zhao, Y.; Li, Y.; Wu, D.H.; Wang, T.; Lu, W.J.; Cai, W.J.; et al. Development and interpretation of a pathomics-based model for the prediction of microsatellite instability in Colorectal Cancer. *Theranostics* **2020**, *10*, 11080. [[CrossRef](#)]
14. Yamashita, R.; Long, J.; Longacre, T.; Peng, L.; Berry, G.; Martin, B.; Higgins, J.; Rubin, D.L.; Shen, J. Deep learning model for the prediction of microsatellite instability in colorectal cancer: A diagnostic study. *Lancet Oncol.* **2021**, *22*, 132–141. [[CrossRef](#)]

15. Echle, A.; Grabsch, H.I.; Quirke, P.; van den Brandt, P.A.; West, N.P.; Hutchins, G.G.; Heij, L.R.; Tan, X.; Richman, S.D.; Krause, J.; et al. Clinical-grade detection of microsatellite instability in colorectal tumors by deep learning. *Gastroenterology* **2020**, *159*, 1406–1416.
16. Ghosh, S.; Santosh, K. Improved Gastrointestinal Screening: Deep Features using Stacked Generalization. In Proceedings of the 2021 IEEE 34th International Symposium on Computer-Based Medical Systems (CBMS), Aveiro, Portugal, 7–9 June 2021; pp. 196–201. [[CrossRef](#)]
17. Lee, S.H.; Song, I.H.; Jang, H.J. Feasibility of deep learning-based fully automated classification of microsatellite instability in tissue slides of colorectal cancer. *Int. J. Cancer* **2021**, *149*, 728–740. [[CrossRef](#)] [[PubMed](#)]
18. Saeed, O.A.; Mann, S.A.; Luchini, C.; Huang, K.; Zhang, S.; Sen, J.D.; Piredda, M.L.; Wang, M.; Baldrige, L.A.; Sperling, R.M.; et al. Evaluating mismatch repair deficiency for solid tumor immunotherapy eligibility: Immunohistochemistry versus microsatellite molecular testing. *Hum. Pathol.* **2021**, *115*, 10–18. [[CrossRef](#)] [[PubMed](#)]
19. Howard, F.M.; Dolezal, J.; Kochanny, S.; Schulte, J.; Chen, H.; Heij, L.; Huo, D.; Nanda, R.; Olopade, O.I.; Kather, J.N.; et al. The Impact of Digital Histopathology Batch Effect on Deep Learning Model Accuracy and Bias. *bioRxiv* **2020**. [[CrossRef](#)]
20. Anghel, A.; Stanislavljevic, M.; Andani, S.; Papandreou, N.; Rüschoff, J.H.; Wild, P.; Gabrani, M.; Pozidis, H. A high-performance system for robust stain normalization of whole-slide images in histopathology. *Front. Med.* **2019**, *6*, 193. [[CrossRef](#)]
21. Zhao, Q.; Adeli, E.; Pohl, K.M. Training confounder-free deep learning models for medical applications. *Nat. Commun.* **2020**, *11*, 1–9. [[CrossRef](#)]
22. Piñol, V.; Andreu, M.; Castells, A.; Payá, A.; Bessa, X.; Jover, R.; of the Spanish Gastroenterological Association, G.O.G. Frequency of hereditary non-polyposis colorectal cancer and other colorectal cancer familial forms in Spain: A multicentre, prospective, nationwide study. *Eur. J. Gastroenterol. Hepatol.* **2004**, *16*, 39–45.
23. Abulí, A.; Bessa, X.; González, J.R.; Ruiz-Ponte, C.; Cáceres, A.; Muñoz, J.; Gonzalo, V.; Balaguer, F.; Fernández-Rozadilla, C.; González, D.; et al. Susceptibility genetic variants associated with colorectal cancer risk correlate with cancer phenotype. *Gastroenterology* **2010**, *139*, 788–796. [[CrossRef](#)] [[PubMed](#)]
24. Macenko, M.; Niethammer, M.; Marron, J.S.; Borland, D.; Woosley, J.T.; Guan, X.; Schmitt, C.; Thomas, N.E. A method for normalizing histology slides for quantitative analysis. In Proceedings of the 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Boston, MA, USA, 28 June–1 July 2009; pp. 1107–1110.
25. Howard, J.; Gugger, S. Fastai. 2018. Available online: <https://github.com/fastai/fastai> (accessed on 1 February 2020).
26. Bankhead, P.; Loughrey, M.B.; Fernández, J.A.; Dombrowski, Y.; McArt, D.G.; Dunne, P.D.; McQuaid, S.; Gray, R.T.; Murray, L.J.; Coleman, H.G.; et al. QuPath: Open source software for digital pathology image analysis. *Sci. Rep.* **2017**, *7*, 1–7.
27. Székely, G.J.; Rizzo, M.L.; Bakirov, N.K. Measuring and testing dependence by correlation of distances. *Ann. Stat.* **2007**, *35*, 2769–2794. [[CrossRef](#)]
28. Lundberg, S.M.; Lee, S.I. A unified approach to interpreting model predictions. In Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 4768–4777.
29. Mercaldo, N.D.; Lau, K.F.; Zhou, X.H. Confidence intervals for predictive values with an emphasis to case–control studies. *Stat. Med.* **2007**, *26*, 2170–2183. [[CrossRef](#)] [[PubMed](#)]