ARTICLE   OPEN

Molecular Diagnostics

Predicting the HER2 status in oesophageal cancer from tissue microarrays using convolutional neural networks

Juan I. Pisula[1], Rabi R. Datta[2], Leandra Börner Valdez[2], Jan-Robert Avemarg[3], Jin-On Jung[2], Patrick Plum[2], Heike Löser[4], Philipp Lohneis[4], Monique Meuschke[3], Daniel Pinto dos Santos[5], Florian Gebauer[2], Alexander Quaas[4], Axel Walch[6], Christiane J. Bruns[2], Kai Lawonn[3], Felix C. Popp[2,8] and Katarzyna Bozek [1,7,8✉]

BACKGROUND: Fast and accurate diagnostics are key for personalised medicine. Particularly in cancer, precise diagnosis is a prerequisite for targeted therapies, which can prolong lives. In this work, we focus on the automatic identification of gastroesophageal adenocarcinoma (GEA) patients that qualify for a personalised therapy targeting epidermal growth factor receptor 2 (HER2). We present a deep-learning method for scoring microscopy images of GEA for the presence of HER2 overexpression.

METHODS: Our method is based on convolutional neural networks (CNNs) trained on a rich dataset of 1602 patient samples and tested on an independent set of 307 patient samples. We additionally verified the CNN's generalisation capabilities with an independent dataset with 653 samples from a separate clinical centre. We incorporated an attention mechanism in the network architecture to identify the tissue regions, which are important for the prediction outcome. Our solution allows for direct automated detection of HER2 in immunohistochemistry-stained tissue slides without the need for manual assessment and additional costly in situ hybridisation (ISH) tests.

RESULTS: We show accuracy of 0.94, precision of 0.97, and recall of 0.95. Importantly, our approach offers accurate predictions in cases that pathologists cannot resolve and that require additional ISH testing. We confirmed our findings in an independent dataset collected in a different clinical centre. The attention-based CNN exploits morphological information in microscopy images and is superior to a predictive model based on the staining intensity only.

CONCLUSIONS: We demonstrate that our approach not only automates an important diagnostic process for GEA patients but also paves the way for the discovery of new morphological features that were previously unknown for GEA pathology.

BACKGROUND

Gastroesophageal adenocarcinoma (GEA) is the seventh most common cancer worldwide, with an increasing number of cases in the western hemisphere. Despite multimodal therapies with neoadjuvant chemotherapy/chemoradiation before surgery, median overall survival does not exceed 4 years [1–5]. Epidermal growth factor receptor 2 (HER2) encodes a transmembrane tyrosine kinase receptor and is present in different tissues, e.g., epithelial cells, mammary gland, and the nervous system. It is also an important cancer biomarker. HER2 activation is associated with angiogenesis and tumorigenesis. Various solid tumours display HER2 overexpression, and targeted HER2 therapy improves their treatment outcomes [6]. Clinical guidelines for GEA recommend adding Trastuzumab—a monoclonal antibody binding to

HER2—to the first-line palliative chemotherapy for HER2-positive cases. HER2 targeting drugs are also currently investigated in the curative therapy for GEA [7].

Accurate testing for the HER2 status is a mandatory prerequisite for the application of targeted therapies. The gold standard for determining the HER2 status is an analysis of the immunohistochemical (IHC) HER2 staining by an experienced pathologist, if necessary followed by an additional in situ hybridisation (ISH). The pathologist examines the immunohistochemistry staining of cancer tissue slides for HER2 and determines the IHC score ranging from 0 to 3. According to expert guidelines [8], the factors determining the score include the staining intensity, the number of connected positive cells, and the cellular location of the staining (Supplemental Table 1). The IHC scores 0 and 1 define patients

[1]Data science of Bioimages Lab, Faculty of Medicine and University Hospital Cologne, Center for Molecular Medicine Cologne (CMMC), University of Cologne, 50931 Cologne, Germany. [2]Department of General, Visceral, Cancer and Transplantation Surgery, University of Cologne, 50937 Cologne, Germany. [3]Faculty of Mathematics and Computer Science, Friedrich Schiller University Jena, 07743 Jena, Germany. [4]Institute of Pathology, University of Cologne, 50937 Cologne, Germany. [5]Department of Radiology, University of Cologne, 50937 Cologne, Germany. [6]Research Unit Analytical Pathology, Helmholtz Zentrum München, 85764 Neuherberg, Germany. [7]Cologne Excellence Cluster on Cellular Stress Responses in Aging-Associated Diseases (CECAD), University of Cologne, Cologne, Germany. [8]These authors contributed equally: Felix C. Popp, Katarzyna Bozek. ✉email: k.bozek@uni-koeln.de

with a negative HER2 status that are not eligible for targeted anti-HER2 therapy. An IHC score of 3 designates a positive HER2 status, and these patients receive Trastuzumab. A score of 2 is equivocal. In this case, an additional in situ hybridisation (ISH) assay resolves the IHC score 2 as a positive or negative HER2 status. However, both manual scoring and additional ISH testing are time-consuming and costly.

Automated IHC quantification can support pathologists and is one of the challenges in digital pathology and Convolutional Neural Network (CNN)-based approaches currently offer the highest accuracy in this task [9]. Tewary and Mukhopadhyay using patch-based labelling created a three-level HER2 classifier with an accuracy of 0.93 [10]. Han et al. combined a patch-level classifier with a second one predicting HER2 score of a whole slide image [11] achieving an accuracy of 0.94. The limitation of these methods is the need for patch-level labelling, which is not typically done in clinical evaluation. Annotations of individual patches are not available in clinical datasets and thus require additional manual work while patch-level predictions require developing aggregation strategies to generate a prediction for the entire slide. Additionally, all of the automated methods to date focus on breast tumours, which have high prevalence and offer several large public datasets. HER2 is however an important biomarker in other cancers, notably oesophageal carcinoma.

Here, we ask whether CNNs can directly predict the HER2 status from IHC-stained tissue sections without additional ISH testing. We investigate which image features the neural network learns to make the prediction—whether it uses only the colour intensity or additional morphological features. We explore a large tissue microarray (TMA) with 1602 digitised images stained for HER2. We use this image dataset as a training set to train two different CNN classification models. We test these models on an independent test dataset of 307 TMA images from an unrelated patient group from the same centre. We also further validate the HER2 status prediction accuracy of our approach on a patient cohort from a different clinical centre. If successful, CNNs could assist pathologists in evaluating IHC stainings and, therefore, save time and expenses related to the ISH analysis.

METHODS
Tumour sample and image preparation
For training the CNNs, we used a multi-spot tissue microarray (TMA) with 165 tumour cases and a single-spot TMA with 428 tumour cases, as described elsewhere [12]. We additionally prepared an independent single-spot TMA with 307 tumour cases as the test dataset. The test set consisted of tumour cases that occurred at a later time point compared to the training set cases. This dataset construction strategy mimics how such a model would be developed and deployed in a clinical routine. Coincidentally, our test set does not include tumour cases with an IHC score of 1. The multi-spot TMA was composed of eight tissue cores (1.2 mm diameter) of each tumour—four cores punched on the tumour margin and four in the tumour centre. To construct the single-spot TMA, we punched one tissue core per patient from the tumour centre. The cores were transferred to TMA receiver blocks. Each TMA block contained 72 tissue cores. Subsequently, we prepared 4 μm-thick sections from the TMA blocks and transferred them to an adhesive-coated slide system (Instrumedics Inc., Hackensack, NJ).

We used a HER2 antibody (Ventana clone 4B5, Roche Diagnostics, Rotkreuz, Switzerland) on the automated Ventana/Roche slide stainer to perform immunohistochemistry (IHC) on the TMA slides. HER2 expression in carcinoma cells was assessed according to staining criteria listed in Supplemental Table 1. Scores 0 and 1 indicated negative HER2 status, and score 3 indicated positive HER2 status. Immunohistochemical expression evaluation was assessed manually by two pathologists (A.Q. and H.L.) according to [13]. Discrepant results, which occurred only in a small number of samples, were resolved by consensus review. Spots with a score of 2 were analysed by fluorescence ISH to resolve the HER2 status. The ISH analysis evaluated the HER2 gene amplification status using the Zytolight SPEC ERBB2/CEN 17 Dual Probe Kit (Zytomed Systems GmbH, Germany)

according to the manufacturer's protocol. A fluorescence microscope (DM5500, Leica, Wetzlar, Germany) with a 63× objective was used for scanning the tumour tissue for amplification hotspots. We counted the signals in randomly chosen areas of homogeneously distributed signals. Twenty tumour cells were evaluated by counting green HER2 and orange centromere-17 (CEN17) signals. The reading strategy followed the recommendations of HER2/CEN17 ratio ≥ 2.0 or HER2 signals ≥ 6.0 for HER2 positive and a HER2/CEN17 ratio <2.0 for HER2-negative samples.

We digitised the slides with a slide scanner (NanoZoomer S360, Hamamatsu Photonics, Japan) with 40-times magnification and used QuPath's [14] TMA dearrayer to slice the digitised slides into individual images (.jpg files, 5468×5468 pixels). After discarding corrupted images, this procedure yielded 1281 images for training, 321 validation, and 307 images for testing. The test set is from the same hospital as the train set but was sampled in a time interval disjoint from and following the time interval when the training dataset was collected. This study design not only reflects potential real life clinical scenarios in which incoming patient data is analysed with a model trained on data collected at an earlier time point, but also it follows the guidelines formulated by Kleppe et al. [15].

To study the capability of the CNNs to generalise, we performed a stringent evaluation of the model performance on an external cohort with 653 samples from a different, geographically separate clinical centre [16]. The same antibody was used to perform the HER2 staining, but the slides showed certain deterioration due to aging. Each image was labelled with the IHC score (0, 1, 2, or 3) and the HER2 status (0 or 1) that was determined by the pathologists or by ISH analysis in equivocal cases. This methodology corresponds to the gold standard, and we used this labelling as ground truth.

Classification models
We implemented a method that allows training neural networks on large images at their original resolution by exploiting weakly supervised Multiple-instance learning (MIL) [17]. In the weakly supervised multiple-instance-learning approach, each slide is considered as a bag of smaller tiles (instances) whose respective individual labels are unknown. To make a bag-level prediction, image tiles are embedded in a low-dimensional vector space, and the embeddings of individual tiles are aggregated to obtain representation of the entire image. This representation is used as input of a bag-level classifier.

For the aggregation of the tile embeddings, we used the attention-based operator proposed by Ilse et al. [18]. It consists of a simple feed-forward network that predicts an attention score for each of the embeddings. These scores indicate how relevant each tile is for the classification outcome, and are used to calculate a weighted sum of the tile representations as the aggregation operation. Weights of a bag sum to one, this way the bag representation is invariant to bag size. Finally, the bag vector representation is used as the input of a feed-forward neural network to perform the final classification.

In this approach, non-overlapping tiles of 224×224 pixels were extracted from each slide, and their embeddings were derived from a ResNet34 model. Empty tiles were discarded beforehand. As in the fully supervised approach, the MIL classifier was trained separately to predict IHC score and HER2 status.

To test the importance of image resolution in prediction we used a ResNet34 architecture [19] for prediction of IHC score and HER2 status. The network was trained as a four class IHC score classifier and separately as a binary classifier of the HER2 status. Given the large resolution of the tissue images (5468×5468 pixels), this approach required scaling them down by 5.34 to the size of 1024×1024 pixels to allow the network to train within our hardware memory limits.

We also constructed a method for predicting IHC score and HER2 status based on the staining intensity of the slides, a feature that is conventionally used by automatic IHC scoring software. This method was constructed to compare how predictive the single feature of staining intensity is compared to the higher level features learned by our CNN models. To extract the IHC staining expression from the images we used colour deconvolution [20]. From the staining channel, non-overlapping tiles of 224×224 pixels were extracted and the average staining intensity was calculated for each tile. The staining intensity of each slide was then calculated as the maximum of the average intensities of its tiles. The proposed slide descriptor was used as input in two logistic regression classifiers to predict IHC score and HER2 status separately. This approach can also be seen as a multiple-instance classification formulation where the feature extracted for each instance
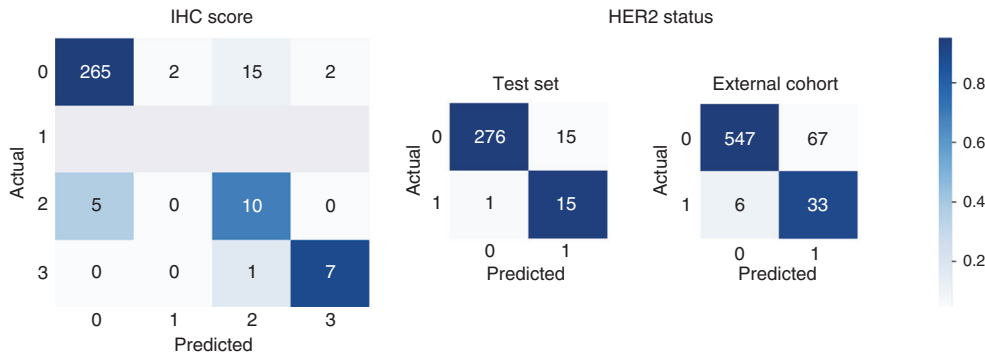
Fig. 1 Confusion matrices of the IHC score and status prediction. Score prediction evaluated on the test set is shown on the left, and HER2 status prediction evaluated on the test set and on the external cohort are shown in the middle and on the right, respectively.

Table 1. Results of the Attention-Based MIL method on the tasks of IHC score prediction and HER2 status prediction.

Task	Balanced acc.	Precision	Recall	F1 score
IHC score prediction	0.8249	0.9470	0.9185	0.9302
HER2 status prediction	0.9429	0.9705	0.9478	0.9551

is its average staining intensity value, and the bag is aggregated using the maximum operator.

Network training

The dataset showed an unbalanced distribution of the IHC score (Supplemental Fig. 1) reflecting the frequency of HER2 expression in the population [21]. To obtain representative training and validation sets, we split images of each IHC score in 80-20 proportions. For the samples with score 2, the 80-20 split was done separately for those with positive status and those with negative status. During training, we performed a weighted sampling of the images of each score such that each of the IHC scores is equally represented during training. We performed random horizontal and vertical flips as data augmentation.

We used Adam optimiser in training [22], with weight decay of 1×10^{-8} and betas of 0.9 and 0.999. The learning rates as well as their schedulers were chosen based on a hyperparameter search. The ResNet classifiers were trained using a learning rate 1×10^{-5}, which was reduced by a factor of 0.1 if the accuracy of the validation set does not improve after 20 epochs of training. The MIL classifier was trained using a learning rate of 5×10^{-9}, decreasing it by a factor of 0.3 if the accuracy of the validation set does not improve after 40 epochs. We used a batch size of 32 in the ResNet classifier and a batch size of only one full resolution image with a bag size depending on the amount of extracted tiles in the MIL classifier.

Our study is compliant with the guidelines summarised by Kleppe et al. [15]. We perform data augmentations, our test set is disjoint in time from the train set, and we demonstrate the method's performance on an external validation set. Our primary analysis was predefined and we report balanced accuracy metrics throughout this study.

Computational work was performed on the CHEOPS high performance computer, on nodes equipped with 4 NVIDIA V100 Volta graphics processing units (GPUs). We used PyTorch (version 1.8.1) [23] for data loading, creating models, and training.

RESULTS

IHC score prediction

First, we implemented a multiple-instance-learning (MIL) [17] method allowing us to make the classification of the images at their highest resolution. Using this technique, the images are split into smaller tiles, encoded into their numeric embeddings and ranked using the attention mechanism as proposed by Ilse et al. [18]. The attention mechanism allows for automatic identification of areas in the image that are important for the predicted score, this way providing means to inspect and interpret the prediction outcomes of the network.

This technique has shown a balanced accuracy of 0.8249, precision of 0.9470 and recall of 0.9185 (Fig. 1: left, Table 1). Given the score imbalance and the lack of samples with an IHC score 1 in the test set, the reported performance metrics were calculated in a balanced manner as an average of the metric of each individual label weighted by their number of samples of that given label. Most notably, the outermost classes 0 and 3 were predicted with the highest accuracy while ~ 33% of score 2 images were incorrectly predicted.

We next examined whether a simpler CNN-based classification approach allows for predicting the IHC score from the TMA images. In order for these images to fit within our hardware constraints, we downsampled them by a factor of 5.34 to a size of 1024×1024 pixels. We trained classification architecture ResNet34 [19] on the rescaled dataset and analysed it on the test set of images adjusted correspondingly. This approach resulted in balanced accuracy of 0.8536, precision of 0.9544 and recall of 0.8859. The almost equal accuracy and precision of this model suggests that relatively large visual details visible at a lower resolution are sufficient for the most accurate prediction.

HER2 status prediction

We next addressed the question whether the HER2 status can be predicted from the IHC-stained images directly, without additional ISH testing. Images in our dataset with IHC score of 0 or 1 are HER2 negative, those with a score of 3 are positive. Those with a score of 2 were additionally resolved using ISH resulting in the following positive/negative HER2 status split: 77/33% in the train set, 53/47% in the test set. Out of 15 IHC score 2 images in the test set, there were eight HER2 positive and seven HER2 negative. The train-validation split was done in such a way that all the score and status combinations are distributed equally in both sets.

The MIL classifier resulted in performance with balanced accuracy of 0.9429, precision of 0.9705 and recall of 0.9478 (Fig. 1 and Table 1). As in the IHC score prediction task, the results were calculated as a weighted average of the individual metrics for class 0 (HER2 negative) and class 1 (HER2 positive) to take account of the class imbalance. Within both the HER2-negative and HER2-positive classes, less than 7% of images were misclassified resulting in balanced precision and recall >0.94. To better understand the errors of the model, we additionally inspected the HER2 status prediction accuracy within images of different IHC

scores (Table 2). With ~27% false-positive and ~7% false-negative predictions, the highest error rate occurred in images with the IHC score of 2. The higher proportion of false positives among the score 2 images could be due to the underrepresentation of samples with this IHC score and negative HER2 status in the training set in the score 2 images. In images with IHC scores 0 and 3, the prediction error was below 4%. The difference in performance between the 4-class and the binary classifiers suggests that the inter-score differences are more subtle than the ones differentiating the two HER2 statuses.

Performance on external cohort

Even if independently, our train and test datasets were collected and prepared within one hospital. To verify how the performance of our model is dependent on the aspects related to the data preparation, we evaluated our models on an independent cohort from a different clinical centre [16]. In particular, we aimed to investigate whether HER2 status prediction is indeed possible using IHC-stained images only. The external cohort included 653

tissue samples belonging to 297 patients with the following IHC score distribution: 416/186/14/37 samples of scores 0/1/2/3 respectively. Out of the score 2 samples, 12 showed a negative HER2 status and 2 samples showed positive HER2 status.

Given the different colour distribution and potential staining quality deterioration due to the sample age, we applied a preprocessing step to these images. We used Macenko's method for stain estimation [24] together with colour deconvolution/ convolution [20] to match the staining to our in-house dataset. The MIL classifier yielded a balanced accuracy of 0.8688, precision of 0.9490 and recall of 0.8908 (Fig. 1). These results support the applicability of our approach in an important clinical context where the distinction of HER2 status is key for further treatment.

Insights into the learning process of the MIL classifier

The ResNet and the MIL classifiers achieved almost identical accuracy on our in-house test set in both the IHC score and the HER2 status prediction. However, the advantage of the more compute-intensive weakly supervised MIL approach is the possibility to inspect the visual features that the network utilises in the classification process. The embeddings and attention scores assigned to individual 224×224 pixel tiles can provide insights into the key visual features used by the MIL approach in the classification.

First, we examined via t-distributed stochastic neighbour embedding (t-SNE) dimensionality reduction method [25] the embeddings of the image tiles in the test set generated by the IHC score prediction network (Fig. 2). In this visualisation, spatial proximity of tiles reflects the similarity of their embeddings. Although the network was trained on the IHC score, it also correctly separates the HER2 status of the parent TMA image. HER2-negative tiles with a

Table 2. Cross-tabulation of true IHC score and predicted HER2 status of the test dataset. '2−' and '2+' scores stand for IHC score 2 and HER2-negative and -positive status, respectively.

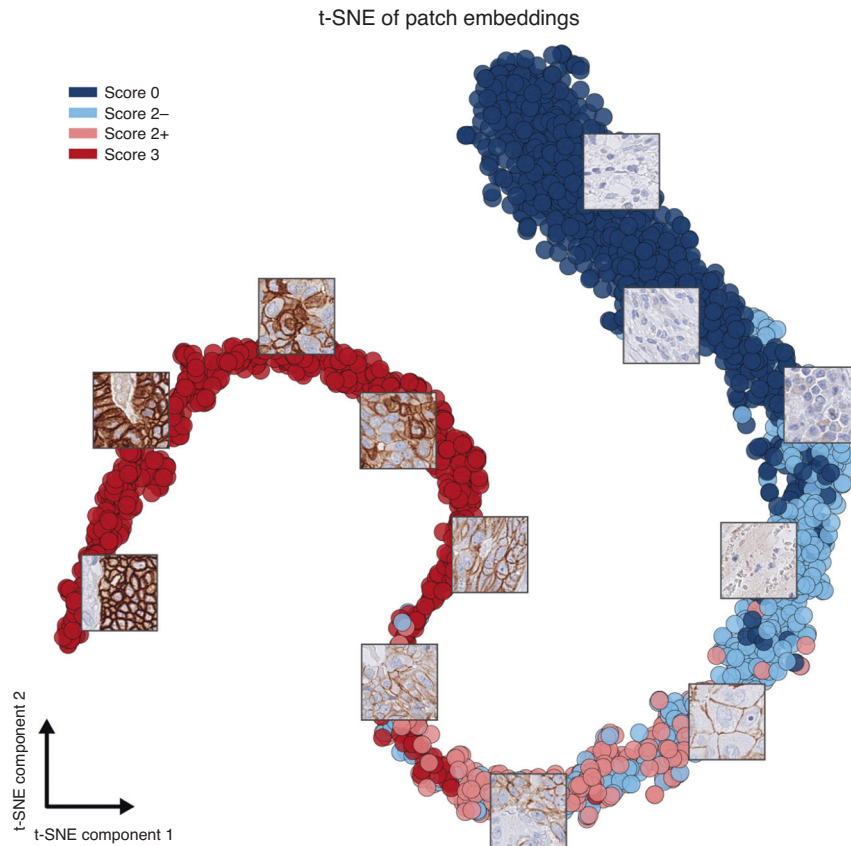|	True IHC score			
Predicted HER2 status	**0**	**2−**	**2+**	**3**
Negative	273	3	1	0
Positive	11	4	7	8

Fig. 2 t-SNE visualisation of tile embeddings produced by the IHC score MIL classifier on the test set images, with the vectors coloured according to the score of their respective slides. Visual similarity of the tiles is reflected in their neural network-derived representations and the embeddings of similar tiles are close in the learned vector space. Coincidentally, there are no TMA images with a score of 1 in the test set because the test set consisted of the consecutive tumour cases that followed the training set cases.
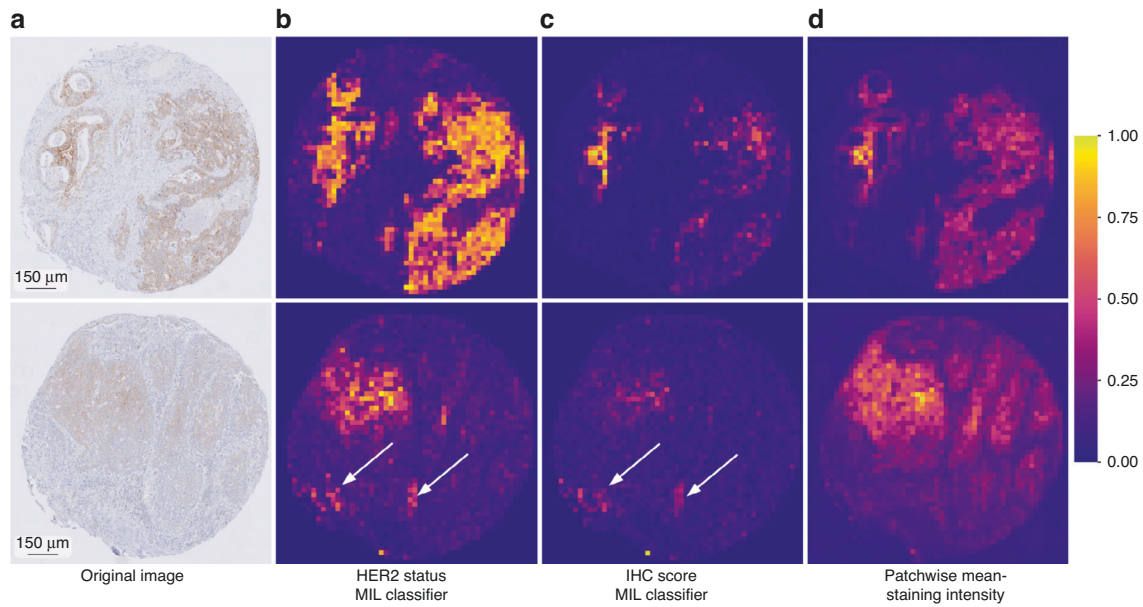
Fig. 3 Heatmap visualisations of the attention value and mean-staining intensity in tiles within the tissue image. The values are normalised to [0, 1]. **a** Slides with IHC score 2 and negative HER2 status. **b** Attention score heatmap of HER2 status MIL classifier. **c** Attention score heatmap of IHC score MIL classifier. **d** Patchwise mean-staining intensity heatmap. White arrows point to locations where the attention values do not match staining intensity.
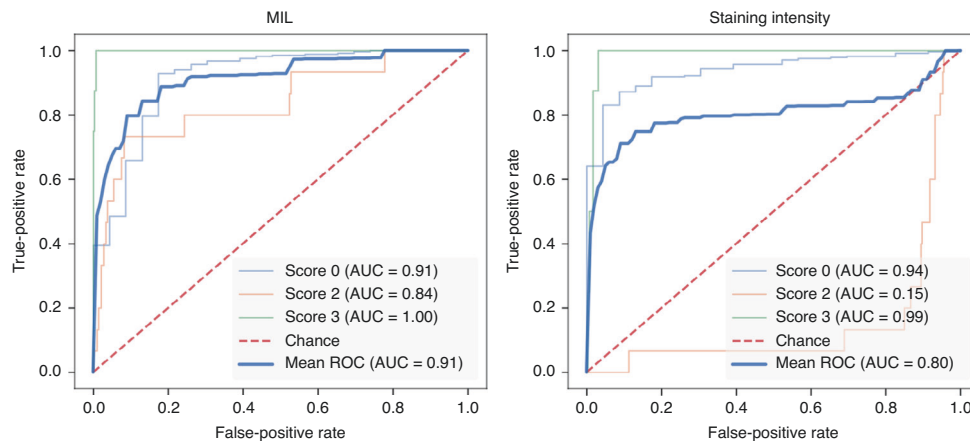

Fig. 4 Per-class ROC curves for the IHC score classifiers, calculated in a "one-vs.-all" fashion of the MIL (left panel) and staining intensity-based classifier (right panel). While both models' performance is similar for images of score 0 and 3, images of score 2 are not possible to correctly recognise based on staining intensity only.

score of 2 (2−) group together with score 0 tiles, and HER2-positive tiles with a score of 2 (2+) group together with score 3 tiles.

Additionally, neighbouring tiles in the t-SNE projection show visual similarity. Most strikingly, tiles grouped together show a similar staining intensity and this intensity gradually changes along the 2D projection of the embeddings. Staining intensity is, however, not the only visual feature determinant of the HER2 scoring, which also takes additional morphological features into account (potentially such as those listed in Supplemental Table 1). We expect these morphological features to also be encoded in the learned vector space.

Next, we inspected the attention values of the MIL classifier and their distribution within the tissue slides. The attention value reflects the importance of a given image tile for the final prediction score and this way provides information on spatial distribution of the visual features in the tissue that the network is exploiting in the prediction. Since the IHC staining is insufficient to resolve the HER2 status if the tissue IHC score is 2, we inspected which visual features are exploited

by the network in resolving the HER2 status of the score 2 tissue slides (Fig. 3). Strikingly, the attention of the MIL classifier for the HER2 status focuses on areas of high staining intensity and corresponds to the mean intensity of the tiles at first sight.

Given the relationship of the embeddings as well as attention value to the staining intensity, we tested the accuracy of a predictive model based on the staining intensity only. Similar to the tiling approach of the MIL classifier, we split the tissue slides in 224×224 pixel tiles and averaged the staining intensity in each of the tiles. We, then, used the maximum of the average intensities across the tiles of an image as the quantitative descriptor of the entire image. We trained two logistic regression models to predict IHC score and HER2 status, respectively. The stain intensity-based model showed a balanced accuracy of 0.6876 in the prediction of the IHC score, markedly lower compared to the MIL classifier with a balanced accuracy of 0.8249. The major difference in performance between these models is in images with an IHC score of 2 (Fig. 4). In the task of predicting the HER2 status, the balanced

accuracy of the staining intensity-based model reached 0.8457 compared to 0.9429 of the MIL classifier.

These results suggest that not only the staining intensity but also additional morphological features are considered by the deep-learning models in the classification. These features are particularly important for correct recognition of images belonging to the intermediate IHC score 2. We indicate examples of such features in Fig. 3 and Supplemental Fig. 2. Even though attention value and staining intensity largely match, the heatmaps in Fig. 3 demonstrate prominent exceptions where features of high attention do not show high staining intensity.

Comparison to existing classifiers
Several computational toolboxes currently allow for training predictive models on whole slide images (WSIs) stained using hematoxylin and eosin (H&E) [26–29]. We compared the results of our approach against CLAM [26], a publicly available pipeline for WSI classification. This pipeline extends the attention-based deep MIL proposed in [18] by including a clustering performed on the embedding space during training, which improves prediction. Similar to our approach, CLAM performs weighted sampling of images to overcome the class imbalance bias. Training and testing CLAM on the same data as our method resulted in balanced accuracy of 0.7166 (precision of 0.9479, recall of 0.7394) in the score prediction task and balanced accuracy of 0.8997 (precision of 0.9611, recall of 0.9218) in the status prediction task, markedly lower compared to our approach.

DISCUSSION
Automated and accurate image-based diagnostics help to accelerate medical treatment and decrease the work burden of the medical personnel. Here, we demonstrate that deep-learning-based prediction of the IHC score (0–3) and the HER2 status (negative or positive) is generally possible with a balanced accuracy of ~0.85 and ~0.94, respectively. Among the scores, IHC score 2 images show the highest proportion of misclassified samples. These score 2 images cannot be unequivocally classified regarding their HER2 status by the pathologists and need further ISH-based evaluation. While it is considered that it is not possible to resolve the HER2 status based on the IHC staining of the IHC score 2 images, our models correctly predict the HER2 status of 73% of these images in our test dataset. Notably, score 2 samples are strongly underrepresented in our datasets. We expect that with more training samples of the underrepresented scores this prediction accuracy will improve.

Several computational toolboxes currently allow for training predictive models on WSIs. These multipurpose pipelines for digital pathology are crucial to the research community because they produce good results, allow for quick insights in the data with an enormous ease of use. Our comparison with an existing, publicly available WSI classification toolbox CLAM [26], suggests however that problem-tailored approaches such as ours offer refined control over parameterisation and data formatting, which allows to achieve higher accuracy and computational efficiency. Dedicated, problem-specific computational solutions might also be easier to further develop into clinical tools.

One of our key findings is that not only staining intensity—conventionally used in automated prediction tools—but also additional morphological properties are taken into account by the neural networks in the classification. We identified multiple images in which the attention maps of the MIL classifier do not match the staining intensity (Fig. 3). Additionally, prediction based on the intensity yields markedly lower accuracy suggesting that the CNN uses morphological features of the image beyond mere staining intensity. This additional information is key for the CNN to correctly predict the equivocal cases with HER2 score 2. Identification of the specific morphological signatures of HER2

not captured by the staining will require pathologists' as well as computational analysis of the high-attention and low stain intensity regions (Supplemental Fig. 2).

Neural networks for quantification of tumour morphology, especially in the H&E stainings, emerge as a novel approach for detecting tumour features invisible to the human eye, such as those corresponding to DNA mutations. Kather et al. predict microsatellite instability in gastrointestinal tumours directly from H&E stainings [30]. Couture et al. predict various breast cancer biomarkers, including the oestrogen receptor status, with an accuracy > 0.75 [31]. The authors suggest the presence of morphological features indicative of the underlying tumour biology in H&E images accessible to deep-learning methods. Lu et al. predict the HER2 status directly from H&E WSIs in breast cancer using a graph representation of the cellular spatial relationship [32] yielding an area under the receiver operator curve (AUROC) of 0.75 on an independent test set.

While inferring information imperceptible to the human eye from H&E stained tumour slides is a powerful approach pushing the boundaries of digital pathology, we use IHC-stained images in our study. Compared to H&E images, IHC stainings directly visualise the molecular HER2 expression and thus present more specific and interpretable data for pathologists. Our approach explores this information to an extent beyond human perception and staining intensity producing an AUROC curve of 0.91 (see Fig. 4). While leaving a clinical decision up to an automated method is not practiced due to its associated ethical questions, our IHC-based MIL approach could readily be used to assist pathologists. The attention maps could point clinicians to the relevant regions in the IHC images and thus save time and manual workload of clinicians.

In this study, our data is in the form of TMA, our approach is however readily applicable to WSIs and expandable to different file formats. Processing optimisations, such as precalculating tile embeddings prior to inference, might be needed if the volume of WSIs exceed the hardware memory limitations. Our results on the external test set suggest that with appropriate image normalisation our model can generalise to other datasets.

Unexpectedly, the classifiers based on low- (1024 × 1024 pixel) and high- (5468 × 5468 pixel) resolution images achieve matched accuracy. Potentially, the lower resolution used in this study is sufficient to encode the key morphological features of the images. This resolution was the highest that still allowed for training ResNet within our hardware memory. Notably decreasing the size of the images further to 512 × 512 pixel size resulted in the decrease of the model balanced accuracy to 0.8200 for the prediction of IHC score. Unlike in this study, WSIs instead of TMAs are used in the diagnostic pathological assessment. The WSI size is several orders of magnitude larger than the images in our dataset, which does not allow for using simple classification architectures such as ResNet and MIL approaches are typically used instead. Our results suggest however that reducing image resolution even 5-fold does not affect the deep-learning model performance, which could accelerate model training and reduce computational costs of models built on WSIs without compromising their accuracy.

Given the class imbalance of our datasets, we report the balanced accuracy and weighted recall, precision and F1 metrics, as the unbalanced and unweighted metrics may be misleading in describing performance of the models. As an example, if unbalanced, the accuracy score of an IHC score classifier that always predicts score 0 would be 0.92 in our dataset, and an analogous HER2 status classifier would achieve accuracy of 0.94. The unbalanced precision (and subsequently, F1) of our HER2 status classifiers would be similarly inaccurate. If we take, for example, the MIL HER2 status classifier, its unbalanced precision score is 0.51, while its false-positive rate is only 0.04. For these reasons we calculate our accuracy metrics in a balanced manner.

We propose that artificial intelligence-based HER2 status evaluation represents a valuable tool to assist clinicians. In particular, the attention map generated by the MIL classifier can aid the pathologists in their daily work by indicating the image areas of high information content for the evaluation. This approach could facilitate and speed up the manual analysis of large tissue images. The IHC score determination network can easily be transferred to any IHC staining other than HER2, further paving the way for digital pathology. We additionally demonstrate the capacity of our method to perform on samples from external clinical centres with similar prediction accuracy. We expect the power and generalisability of our deep-learning model to increase with larger, multi-centre datasets.

Finally, the high performance of our models in predicting the HER2 status of score 2 samples for which the status is considered as unresolvable based on the IHC staining, suggests that there exist visual features predictive of the HER2 status in these images. While identification of these features would require more IHC score 2 image data than available in our dataset, we expect that further deployment of the MIL models might lead to the discovery of novel morphological signatures improving image-based diagnostics.

CONCLUSION
We demonstrate that it is possible to automatically predict HER2 overexpression directly from IHC-stained images of gastroesophageal cancer tissue, an important diagnostic process in the treatment of GEA patients. CNNs not only replicate the IHC scoring system used by pathologists, but can directly predict HER2 status in cases where it is considered not possible to resolve this condition by IHC staining alone.

Interestingly, staining intensity is not the only predictive feature for HER2 overexpression in the IHC images. Deep-learning algorithms can capture complex molecular features like the HER2 status from the tissue morphology. The attention map of the MIL classifier identifies key morphological features beyond staining intensity that might be important indicators to assess individual tumour biology.

We conclude that deep-learning-based image analysis represents a valuable tool both for the development of useful digital pathology applications and the discovery of visual features and patterns previously unknown to traditional pathology.

DATA AVAILABILITY
The data that supports the findings of this study is publicly available in https://zenodo.org/record/7031868 [33].

CODE AVAILABILITY
We provide our code with explanatory notebooks under https://github.com/bozeklab/HER2-overexpression.

REFERENCES
1. Dai T, Shah MA. Chemoradiation in oesophageal cancer. Best Pract Res Clin Gastroenterol. 2015;29:193–209.
2. van Hagen P, Hulshof MC, Van Lanschot JJ, Steyerberg EW, Henegouwen MV, Wijnhoven BP, et al. Preoperative chemoradiotherapy for esophageal or junctional cancer. N Engl J Med. 2012;366:2074–84.
3. Xi M, Hallemeier CL, Merrell KW, Liao Z, Murphy MA, Ho L, et al. Recurrence risk stratification after preoperative chemoradiation of esophageal adenocarcinoma. Ann Surg. 2018;268:289–95.
4. Noordman BJ, Verdam MG, Lagarde SM, Hulshof MC, Hagen PV, van Berge Henegouwen MI, et al. Effect of neoadjuvant chemoradiotherapy on health-related quality of life in esophageal or junctional cancer: results from the randomized CROSS trial. J Clin Oncol. 2018;36:268–75.
5. Shapiro J, Van Lanschot JJ, Hulshof MC, van Hagen P, van Berge Henegouwen MI, Wijnhoven BP, et al. Neoadjuvant chemoradiotherapy plus surgery versus surgery alone for oesophageal or junctional cancer (CROSS): long-term results of a randomised controlled trial. Lancet Oncol. 2015;16:1090–8.
6. Oh DY, Bang YJ. HER2-targeted therapies—a role beyond breast cancer. Nat Rev Clin Oncol. 2020;17:33–48.
7. Wagner AD, Grabsch HI, Mauer M, Marreaud S, Caballero C, Thuss-Patience P, et al. EORTC-1203-GITCG-the "INNOVATION"-trial: Effect of chemotherapy alone versus chemotherapy plus trastuzumab, versus chemotherapy plus trastuzumab plus pertuzumab, in the perioperative treatment of HER2 positive, gastric and gastroesophageal junction adenocarcinoma on pathologic response rate: a randomized phase II-intergroup trial of the EORTC-Gastrointestinal Tract Cancer Group, Korean Cancer Study Group and Dutch Upper GI-Cancer group. BMC Cancer. 2019;19:1–9.
8. Nie J, Lin B, Zhou M, Wu L, Zheng T. Role of ferroptosis in hepatocellular carcinoma. J Cancer Res Clin Oncol. 2018;144:2329–37.
9. Qaiser T, Mukherjee A, Reddy Pb C, Munugoti SD, Tallam V, Pitkäaho T, et al. HER2 challenge contest: a detailed assessment of automated HER2 scoring algorithms in whole slide images of breast cancer tissues. Histopathology. 2018;72:227–38.
10. Tewary S, Mukhopadhyay S. HER2 molecular marker scoring using transfer learning and decision level fusion. J Digit Imaging. 2021;34:667–77.
11. Han Z, Lan J, Wang T, Hu Z, Huang Y, Deng Y, et al. A deep learning quantification algorithm for HER2 scoring of gastric cancer. Front Neurosci. 2022;16:877229.
12. Plum PS, Gebauer F, Krämer M, Alakus H, Berlth F, Chon SH, et al. HER2/neu (ERBB2) expression and gene amplification correlates with better survival in esophageal adenocarcinoma. BMC Cancer. 2019;19:1–9.
13. Lordick F, Al-Batran SE, Dietel M, Gaiser T, Hofheinz RD, Kirchner T, et al. HER2 testing in gastric cancer: results of a German expert meeting. J cancer Res Clin Oncol. 2017;143:835–41.
14. Bankhead P, Loughrey MB, Fernández JA, Dombrowski Y, McArt DG, Dunne PD, et al. QuPath: open source software for digital pathology image analysis. Sci Rep. 2017;7:1–7.
15. Kleppe A, Skrede OJ, De Raedt S, Liestøl K, Kerr DJ, Danielsen HE. Designing deep learning studies in cancer diagnostics. Nat Rev Cancer. 2021;21:199–211.
16. Langer R, Rauser S, Feith M, Nährig JM, Feuchtinger A, Friess H, et al. Assessment of ErbB2 (Her2) in oesophageal adenocarcinomas: summary of a revised immunohistochemical evaluation system, bright field double in situ hybridisation and fluorescence in situ hybridisation. Mod Pathol. 2011;24:908–16.
17. Dietterich TG, Lathrop RH, Lozano-Pérez T. Solving the multiple instance problem with axis-parallel rectangles. Artif Intell. 1997;89:31–71.
18. Ilse M, Tomczak J, Welling M. Attention-based deep multiple instance learning. In: International conference on machine learning. Jul 3. PMLR; 2018. pp. 2127–36.
19. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. IEEE:2016. pp. 770–78.
20. Ruifrok AC, Johnston DA. Quantification of histochemical staining by color deconvolution. Anal Quant Cytol Histol. 2001;23:291–9.
21. Koopman T, Smits MM, Louwen M, Hage M, Boot H, Imholz AL. HER2 positivity in gastric and esophageal adenocarcinoma: clinicopathological analysis and comparison. J Cancer Res Clin Oncol. 2015;141:1343–51.
22. Kingma DP, Ba J. Adam: a method for stochastic optimization. https://arxiv.org/abs/1412.6980. 2014.
23. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems. https://doi.org/10.48550/arXiv.1912.01703. 2019;32.
24. Macenko M, Niethammer M, Marron JS, Borland D, Woosley JT, Guan X, et al. A method for normalizing histology slides for quantitative analysis. In: 2009 IEEE international symposium on biomedical imaging: from nano to macro. IEEE:2009. pp. 1107–10.
25. Van der Maaten L, Hinton G. Visualizing data using t-SNE. J Mach Learn Res. 2008;9:2579–605.
26. Lu MY, Williamson DF, Chen TY, Chen RJ, Barbieri M, Mahmood F. Data-efficient and weakly supervised computational pathology on whole-slide images. Nat Biomed Eng. 2021;5:555–70.
27. van Treeck M, Cifci D, Laleh NG, Saldanha OL, Loeffler CM, Hewitt KJ, et al. DeepMed: a unified, modular pipeline for end-to-end deep learning in computational pathology. https://doi.org/10.1101/2021.12.19.473344. 2021.
28. Dolezal J, Kochanny S, Howard F, Slideflow: a unified deep learning pipeline for digital histology (1.1.0). Zenodo. https://doi.org/10.5281/zenodo.6465196. 2022.
29. Pocock J, Graham S, Vu QD, Jahanifar M, Deshpande S, Hadjigeorghiou G, et al. TIAToolbox: an end-to-end toolbox for advanced tissue image analytics. Commun Med (Lond). 2022;2:120.

30. Kather JN, Pearson AT, Halama N, Jäger D, Krause J, Loosen SH, et al. Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. Nat Med. 2019;25:1054–6.

31. Couture HD, Williams LA, Geradts J, Nyante SJ, Butler EN, Marron JS, et al. Image analysis with deep learning to predict breast cancer grade, ER status, histologic subtype, and intrinsic subtype. NPJ Breast Cancer. 2018;4:30.

32. Lu W, Toss M, Dawood M, Rakha E, Rajpoot N, Minhas F. SlideGraph+: Whole slide image level graphs to predict HER2 status in breast cancer. Med Image Anal. 2022;80:102486.

33. Pisula JI, Datta RR, Boerner-Valdez L, Jung JO, Plum P, Loeser H, et al. HER2 overexpression in gastroesophageal adenocarcinoma from immunohistochemstry imaging (0.1). Zenodo. https://doi.org/10.5281/zenodo.7031868. 2022.

AUTHOR CONTRIBUTIONS
JIP performed data analysis, paper writing and software development. RRD, LBV, JJ, PP, PL, MM, DPS and KL did the data acquisition. JRA contributed to software development. HL, FG, AQ performed data acquisition and analysis. CJB and KL revised the paper. AW provided the external cohort. FCP and KB designed the study and performed data analysis and writing of the paper.

COMPETING INTERESTS
The authors declare no competing interests.

CONSENT FOR PUBLICATION
Not applicable.

ADDITIONAL INFORMATION
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41416-023-02143-y.

Correspondence and requests for materials should be addressed to Katarzyna Bozek.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.