



Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

BACH: Grand challenge on breast cancer histology images[☆]



Guilherme Aresta^{a,b,*}, Teresa Araújo^{a,b,*}, Scotty Kwok^c, Sai Saketh Chennamsetty^d, Mohammed Safwan^e, Varghese Alex^f, Bahram Marami^g, Marcel Prastawa^g, Monica Chan^g, Michael Donovan^g, Gerardo Fernandez^g, Jack Zeineh^g, Matthias Kohl^h, Christoph Walzⁱ, Florian Ludwig^h, Stefan Braunewell^h, Maximilian Baust^h, Quoc Dang Vu^j, Minh Nguyen Nhat To^j, Eal Kim^j, Jin Tae Kwak^j, Sameh Galal^k, Veronica Sanchez-Freire^k, Nadia Brancati^l, Maria Frucci^l, Daniel Riccio^{l,m}, Yaqi Wangⁿ, Lingling Sun^{n,o}, Kaiqiang Maⁿ, Jiannan Fangⁿ, Ismael Kone^p, Lahsen Boulmane^p, Aurélio Campilho^{a,b}, Catarina Eloy^{q,r,s}, António Polónia^{q,r,s}, Paulo Aguiar^{s,t}

^aINESC TEC - Institute for Systems and Computer Engineering, Technology and Science, Porto 4200-465, Portugal

^bFaculty of Engineering of University of Porto, Porto 4200-465, Portugal

^cSeek AI Limited, Hong Kong, China

^dBangalore, India

^eGurgaon, India

^fChennai, India

^gThe Center for Computational and Systems Pathology, Department of Pathology, Icahn School of Medicine at Mount Sinai and The Mount Sinai Hospital, New York, USA

^hKonica Minolta Laboratory Europe, Munich, Germany

ⁱInstitute of Pathology, Faculty of Medicine, LMU Munich, Munich, Germany

^jDepartment of Computer Science and Engineering, Sejong University, Seoul 05006, Korea

^kChicago, IL, USA

^lInstitute for High Performance Computing and Networking, National Research Council of Italy (ICAR-CNR), Naples, Italy

^mUniversity of Naples "Federico II", Naples, Italy

ⁿKey Laboratory of RF Circuits and Systems, Ministry of Education, Hangzhou Dianzi University, Hangzhou 310018, China

^oZhejiang Provincial Laboratory of Integrated Circuits Design, Hangzhou Dianzi University, Hangzhou 310018, China

^p2MIA Research Group, LEM2A Lab, Faculté des Sciences, Université Moulay Ismail, Meknes, Morocco

^qLaboratório de Anatomia Patológica, Ipatimup Diagnósticos, Rua Júlio Amaral de Carvalho, Porto 45, 4200-135, Portugal

^rFaculdade de Medicina, Universidade do Porto, Alameda Prof Hernâni Monteiro, Porto 4200-319, Portugal

^sInstituto de Investigação e Inovação em Saúde (i3S), Universidade do Porto, Rua Alfredo Allen, 208, Porto 4200-135, Portugal

^tInstituto de Engenharia Biomédica (INEB), Universidade do Porto, Rua Alfredo Allen, 208, Porto 4200-135, Portugal

ARTICLE INFO

Article history:

Received 4 October 2018

Revised 28 May 2019

Accepted 29 May 2019

Available online 31 May 2019

ABSTRACT

Breast cancer is the most common invasive cancer in women, affecting more than 10% of women worldwide. Microscopic analysis of a biopsy remains one of the most important methods to diagnose the type of breast cancer. This requires specialized analysis by pathologists, in a task that i) is highly time- and cost-consuming and ii) often leads to nonconsensual results. The relevance and potential of automatic classification algorithms using hematoxylin-eosin stained histopathological images has already been demonstrated, but the reported results are still sub-optimal for clinical use. With the goal of advancing the state-of-the-art in automatic classification, the Grand Challenge on BreAst Cancer Histology images (BACH) was organized in conjunction with the 15th International Conference on Image Analysis and Recognition (ICIAR 2018). BACH aimed at the classification and localization of clinically relevant

[☆] ©2018. Licensed under the Creative Commons CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

* Corresponding authors. These authors contributed equally to this work.

E-mail addresses: guilherme.m.aresta@inesctec.pt (G. Aresta), tfaraujo@inesctec.pt (T. Araújo), apolonia@ipatimup.pt (A. Polónia), pauloaguiar@ineb.up.pt (P. Aguiar).

Keywords:

Breast cancer
 Histology
 Digital pathology
 Challenge
 Comparative study
 Deep learning

histopathological classes in microscopy and whole-slide images from a large annotated dataset, specifically compiled and made publicly available for the challenge. Following a positive response from the scientific community, a total of 64 submissions, out of 677 registrations, effectively entered the competition. The submitted algorithms improved the state-of-the-art in automatic classification of breast cancer with microscopy images to an accuracy of 87%. Convolutional neuronal networks were the most successful methodology in the BACH challenge. Detailed analysis of the collective results allowed the identification of remaining challenges in the field and recommendations for future developments. The BACH dataset remains publicly available as to promote further improvements to the field of automatic classification in digital pathology.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Breast cancer is one of most common cancer-related death causes in women of all age (Siegel et al., 2017), but early diagnosis and treatment can significantly prevent the disease's progression and reduce its morbidity rates (Smith et al., 2005). Because of this, women are recommended to do self check-ups via palpation and regular screenings via ultrasound or mammography; if an abnormality is found, a breast tissue biopsy is performed (American Cancer Society, 2015). Usually, the collected tissue sample is stained with hematoxylin and eosin (H&E), which allows to distinguish the nuclei from the parenchyma, and is observed via an optic microscope. Complementarily, these samples can also be scanned to giga-pixel size images, referred as whole-slide image (WSI), for posterior digital processing. During assessment, pathologists search for signs of cancer on microscopic portions of the tissue by analyzing its histological properties. This procedure allows to distinguish malignant regions from non-malignant (benign) tissue, which present changes in normal structures of breast parenchyma that are not directly related with progression to malignancy. These malignant lesions can be further classified as *in situ* carcinoma, where the cancerous cells are restrained inside the mammary ductal-lobular system, or invasive if the cancer cells are spread beyond the ducts. Due to the importance of correct diagnosis in patient management the search for precise, robust and automated systems has increased. The differentiation of breast samples into normal, benign and malignant (either *in situ* or invasive) brings relevant changes in the treatment of the patients making the accurate diagnosis essential. For instance, benign lesions can usually be followed clinically without the need for surgery, but malignancy almost always require surgery with or without the addition of chemotherapy.

The analysis of breast cancer WSIs is non-trivial due to the large amount of data to visualize and the complexity of the task (Elmore et al., 2015). On this setting, computer-aided diagnosis (CAD) systems can alleviate the procedure by providing a complementary and objective assessment to the pathologist. Despite the high performance of these systems for the binary classification (healthy vs malignant) of microscopy (Kowal et al., 2013; Filipczuk et al., 2013; George et al., 2014; Belsare et al., 2015) and whole-slide images (Cruz-Roa et al., 2014; 2018), the previously referred standard clinical classification procedure has now only started to be explored (Araújo et al., 2017; Fondón et al., 2018; Han et al., 2017; Bejnordi et al., 2017).

1.1. Related work

Automatic methods for breast cancer assessment in histology images can be divided according to the type of image in study (namely microscopy images and WSI) and number of classes, *i.e.*, abnormality types, they consider.

1.1.1. Microscopy images

The classification of breast histology microscopy images as benign-malignant for referral purposes is a vastly addressed topic. Over the past decade, these methods have focused on the extraction of nuclei features, which requires the detection of these regions-of-interest. For example, nuclei have been segmented via color-based clustering (Kowal et al., 2013) or by nuclei candidate detection using the circular Hough transform, followed by feature-based candidate reduction and refinement via watersheds (George et al., 2014). These segmentations allow to extract features, usually related to morphology, topology and texture. The computed features can then be used for training one or more classifiers and allow to achieve accuracies of 84–93% (Kowal et al., 2013) and 72–97% (George et al., 2014).

A viable alternative to the design and extraction of hand-crafted features is to use deep learning approaches, namely convolutional neural networks (CNNs), since these allow to significantly reduce the need for field-knowledge while achieving similar or better results. For instance, (Spanhol et al., 2016b) used CNNs to classify patches of microscopy images, and combined the predictions into an image label through sum, product and maximum rules. The method was evaluated on the BreakHis dataset (Spanhol et al., 2016a), which contains images of different magnifications, and achieved an accuracy of 84% for a $200 \times$ magnification.

The more complex 3-class problem of considering normal tissue, *in situ* carcinoma and invasive carcinoma has also been addressed by the scientific community. Due to the increased complexity of the task, using nuclei-related features is usually not sufficient to achieve a reasonable classification performance. Namely, distinguishing *in situ* and invasive carcinomas requires assessing both the nuclei and their organization on the tissue. For instance, Zhang (2011) used a cascade classification approach, where features based on the curvelet transform and local binary patterns were randomly chosen as input to a set of parallel support-vector machines (SVMs). The images where no agreement was found were analysed by a set of NNs using another random feature set, resulting in an accuracy of 97%.

Despite the successes for 2-class and 3-class classifications, few works have addressed the 4 classification problem (normal tissue, benign lesion, *in situ* and invasive carcinoma) of histology images. Recently Fondón et al. (2018) proposed a handcrafted feature-based approach, considered three major sets of features, related with the nuclei, color regions and textures accounting for local and global image properties, which were then used for training a SVM. By their turn, Araújo et al. (2017) proposed a CNN-based approach, training a VGG-like network using patches extracted from the histology images. In the design of the network the authors had in consideration the effective receptive fields at each network layer in order to ensure that information is captured at different scales, allowing that both nuclei organization and the overall tissue structure could be considered. The features extracted by the CNN were then used to train a SVM, and majority voting was

used for obtaining the final image label from the individual patch classifications. These methods have been developed in the context of the Bioimaging 2015 challenge¹. Fondón et al. (2018) and Araújo et al. (2017) have achieved accuracies in the 4-class problem of around 68% and 78%, respectively, on a test set of 36 images from which approximately a half correspond to extremely hard cases to classify, accordingly to two specialists.

1.1.2. Whole-slide image analysis

The recent advances on the acquisition systems have enabled the digitization of entire slides, avoiding loss of biopsy tissue and providing extra context for pathology assessment by the medical experts. However, automatic analysis of these images is challenging due to their giga-pixel size and wide-variety of local tissue behavior. Because of this, supervised WSI analysis has been mainly performed by assessing patches of different magnifications. For instance, the detection of invasive cancer regions can be performed by training a CNN with small patches and predicting over an entire slide. A properly trained CNN can achieve balanced accuracies (the average of specificity and sensitivity) of 84%, outperforming handcrafted-feature approaches by more than 5% (Cruz-Roa et al., 2014; 2017).

The majority of the solutions applied to WSIs are computationally expensive and deal only with small regions of interest (ROIs) and not the complete WSI, since this would require estimating an incredibly high number of parameters. Thus, efforts have been made for developing methods which can deal with these large sized images without the need for ROI selection or extremely high computational power while yielding a good performance. For instance, Cruz-Roa et al. (2018) proposed the combination of CNNs and an adaptive sampling method which relies on the quasi-Monte Carlo sampling and a gradient-based adaptive strategy aiming at focusing sampling on areas of the image of higher uncertainty. The method was evaluated on 195 studies, achieving a Dice coefficient of 76% and yielded comparable results to those from a dense sampling while greatly increasing the computational efficiency (1500 × faster).

Similarly to microscopy images, the scientific community is now starting to explore multi-class classification on WSI, mainly using deep learning approaches. In (Bejnordi et al., 2017), a context aware stack of 2 CNNs was used for detecting normal/benign tissue as well as *in situ* and invasive carcinomas. The first CNN was trained for classifying small high-resolution patches of the WSIs, thus learning cellular-level characteristics. This fully-convolutional model was then used for predicting a set of feature maps from patches of higher size, which serve as input for a second CNN. This scheme allows to integrate both local and global features related to tissue organization, achieving accuracies of 82% and a Cohen's kappa value of 0.7.

By its turn, Gecer et al. (2018) considered 5 classes for the detection and classification of cancer in WSIs: non proliferative changes only, proliferative changes, atypical ductal hyperplasia, *in situ* and invasive carcinoma. The classification was performed by combining the prediction of two steps. First, four CNNs were used sequentially to detect ROIs in a multiscale fashion. Specifically, the output of each CNN is a set of ROIs for the next, thus increasing the magnification of the analyzed region. Then, the output of the fourth model is classified by a CNN trained with high magnification patches. The outputs of the classification and ROI-proposal CNNs are then combined via majority voting, allowing to obtain a slide-level accuracy of 55%, which revealed no statistical difference from 45 pathologists' performance.

1.2. Challenges

Challenges are known to enable advances on the medical image analysis field by promoting the participation of multiple researchers of different backgrounds on a competitive, but scientifically constructive, setting. Over the past years, the scientific community has been promoting challenges on different imaging modalities and topics. Related to breast cancer, CAMELYON² is a two-edition challenge aimed at cancer metastases detection on WSI of lymph node sections.

To further promote and complement the research on the breast cancer image analysis field, the Grand Challenge on Breast Cancer Histology images (BACH) was organized as part of the ICIAR 2018 conference (15th International Conference on Image Analysis and Recognition)³. BACH is a biomedical image challenge built on top of the Bioimaging 2015 challenge, with a much larger dataset of H&E stained microscopy images for classification and a new set of WSI breast cancer tissue images for segmentation. Specifically, the participants of BACH were asked to predict the type of these tissue samples as 1) Normal, 2) Benign 3) *In situ* carcinoma and 4) Invasive carcinoma, with the goal of providing pathologists a tool to reduce the diagnosis workload. The rest of the paper is organized as follows. Section 2 details the challenge in terms of organization, dataset and participant's evaluation. Then, Section 3 describes the approaches of the best performing methods, and the corresponding performance is provided on Section 4. Finally, Section 6 summarizes the findings of this study.

2. Challenge description

2.1. Organization

The BACH challenge was organized into different stages, providing a well structured workflow to potentiate the success of the initiative (Fig. 1). The challenge was hosted on *Grand Challenge*⁴, which allowed for an easy platform set-up. At the time of this writing, *Grand Challenge* accounts for more than 12,000 registered users and, alongside Kaggle⁵, it is one of the preferred platforms for medical imaging-related challenges. The BACH was also announced via the *Sci-diku-imageworld* mailing list⁶. Participants were asked to register on the *Grand Challenge* to access most of the contents of the BACH webpage. All registrations were manually validated by the organization to minimize spam and anonymous participation. Once accepted, participants could download the data by filling a form asking for their name, institution and e-mail address. Once the form was submitted, an e-mail containing an unique set of credentials (username, password) and the dataset download link was automatically sent to the provided e-mail address. This allowed the organization to better limit the dataset access to non-participants as well as collect a list of the institutions/companies interested in the challenge.

BACH was divided in two parts, **A** and **B**. Part **A** consisted in automatically classifying H&E stained breast histology microscopy images in four classes: 1) Normal, 2) Benign, 3) *In situ* carcinoma and 4) Invasive carcinoma. Part **B** consisted in performing pixel-wise labeling of whole-slide breast histology images in the same four classes. Participants were allowed to participate on a single part of the challenge. Also, to promote participation (thus more competition and higher quality of the methods), the ICIAR 2018 conference sponsored the challenge by awarding monetary prizes

² <https://camelyon17.grand-challenge.org/>.

³ <https://www.aimiconf.org/iciar18/>.

⁴ <https://grand-challenge.org/>.

⁵ <https://www.kaggle.com/>.

⁶ <https://list.ku.dk/listinfo/sci-diku-imageworld>.

¹ http://www.bioimaging2015.ineb.up.pt/challenge_overview.html.

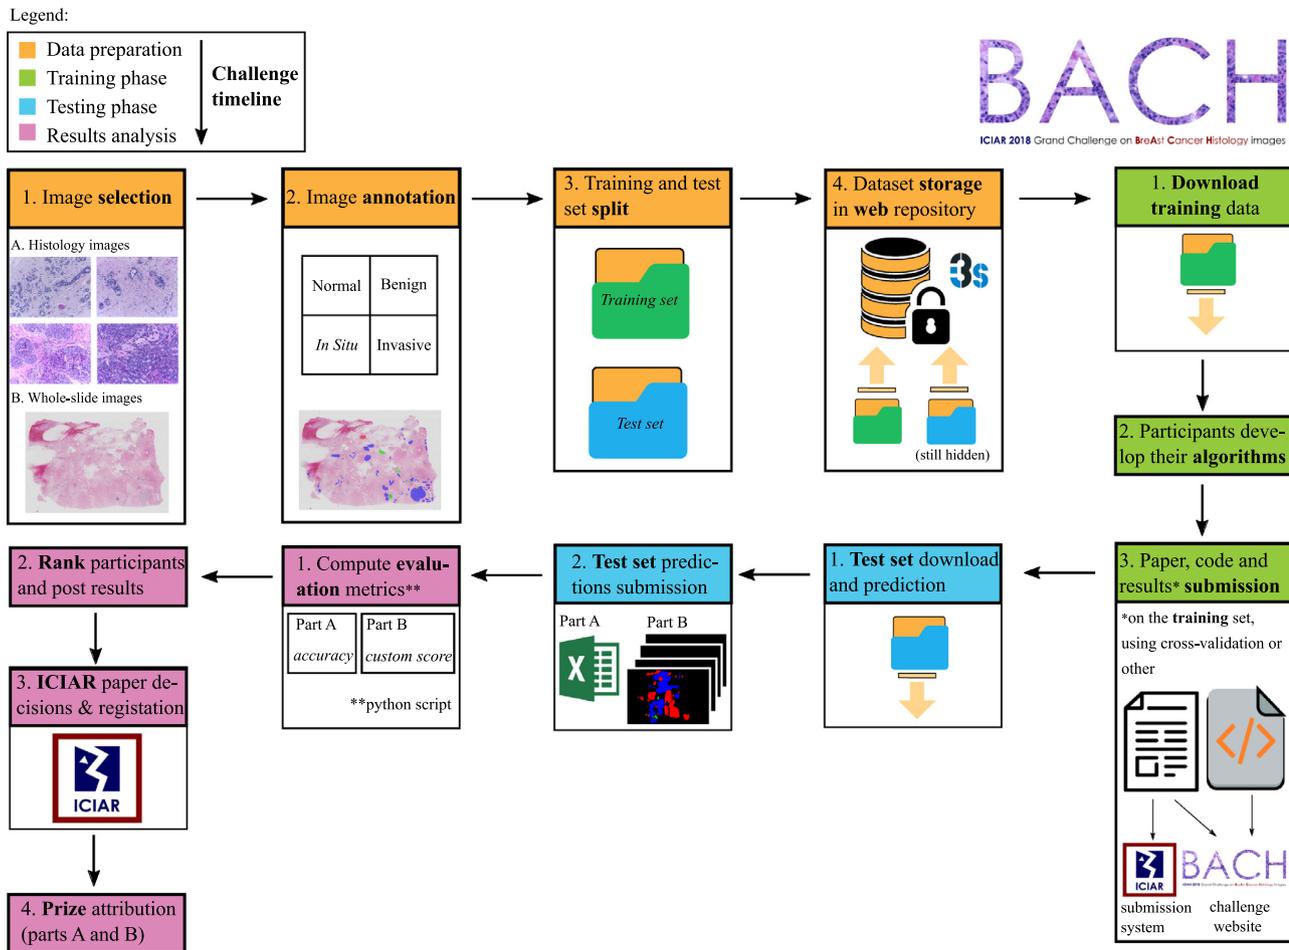


Fig. 1. Workflow of the BACH challenge.

to the first and second best performing methods, for both challenge parts. The prize awarding was contingent of the acceptance and presentation of the methodology at the ICIAR 2018 conference.

The BACH website⁷ was first made publicly available on the 1st November 2017 with the release of the labeled training set. The registered participants had up to 1st February 2018 (4 months) to submit the source code of their methods and a paper describing their approach. To promote the dissemination of the methods, participants were also required to submit their paper to the ICIAR conference. The test set was released on the 5th February 2018 and submissions were open for a week. Results were announced a month after, at the 12th March 2018.

2.2. Datasets

The BACH challenge made available two labeled training datasets for the registered participants. The first dataset is composed of microscopy images annotated image-wise by two expert pathologists from the Institute of Molecular Pathology and Immunology of the University of Porto (IPATIMUP) and from the Institute for Research and Innovation in Health (i3S). The second dataset contains pixel-wise annotated and non-annotated WSI images. For the WSI, annotations were performed by a pathologist and revised by a second expert. The training data is made publicly available from the BACH challenge website: <https://iciar2018-challenge.grand-challenge.org/>

2.2.1. Microscopy images dataset

The microscopy dataset is composed of 400 training and 100 test images, with the four classes equally represented (see Fig. 2). All images were acquired in 2014, 2015 and 2017 using a Leica DM 2000 LED microscope and a Leica ICC50 HD camera and all patients are from the Porto and Castelo Branco regions (Portugal). Cases are from Ipatimup Diagnostics and come from three different hospitals (Hospital CUF Porto, Centro Hospitalar do Tâmega e Sousa and Centro Hospitalar Cova da Beira). The annotation was performed by two medical experts. Images where there was disagreement between the Normal and Benign classes were discarded. The remaining doubtful cases were confirmed via immunohistochemical analysis. The provided images are on RGB .tiff format and have a size of 2048×1536 pixels and a pixel scale of $0.42 \mu\text{m} \times 0.42 \mu\text{m}$. The labels of the images were provided in .csv format. Participants were provided with a partial patient-wise distribution of the images hloof the training set. The test data was collected from a completely different set of patients, ensuring a fairer evaluation of the methods. Note that the training set is an extension of the one used for developing the approach in Araújo et al. (2017).

2.2.2. Whole-slide images dataset

Whole-slide images (WSI) are high resolution images containing the entire sampled tissue. Because of that, each WSI can have multiple pathological regions. The BACH's Part B dataset is composed of 30 WSI for training and 10 WSI for algorithm testing. Specifically for training, the organization provided 10 pixel-wise annotated regions for the Benign, *In situ* carcinoma and Invasive carci-

⁷ <https://iciar2018-challenge.grand-challenge.org/>.

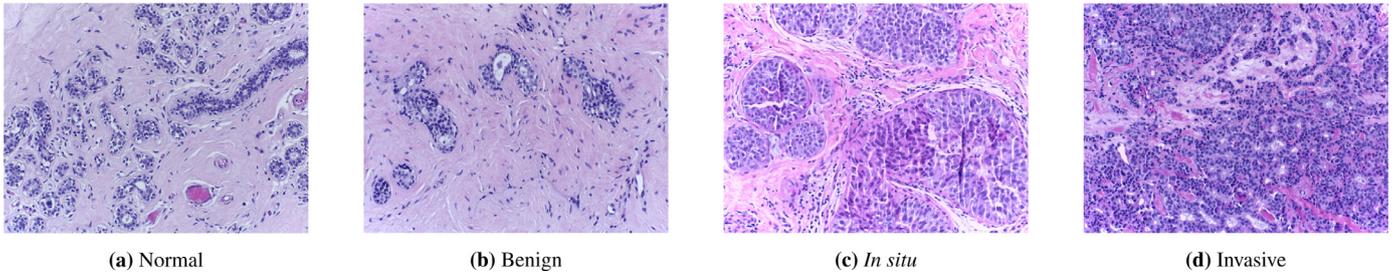


Fig. 2. Examples of microscopy images from the BACH dataset.

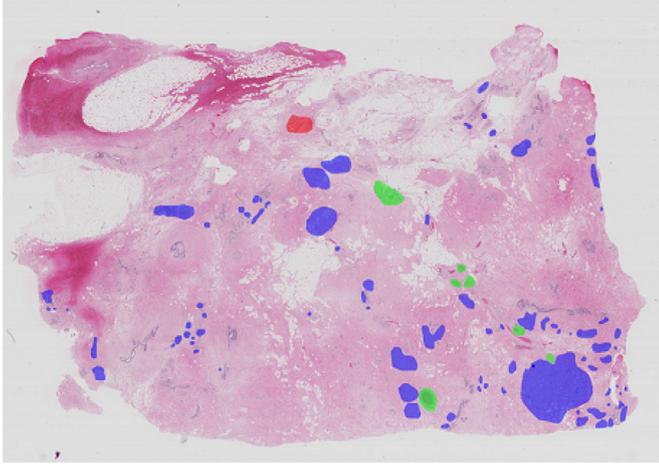


Fig. 3. Example of a pixel-wise annotated whole-slide image from the training set. ■ benign; ■ *in situ*; ■ invasive.

Table 1
Relative distribution (%) of the labels for the training and test sets of Part B.

	Benign	<i>In situ</i>	Invasive
Train	9	3	88
Test	31	6	63

noma classes and 20 potentially pathological WSIs that were not annotated by the experts. The provided annotations aim at identifying regions of interest for the diagnosis on the lowest magnification setting and thus may include non-tissue and normal tissue regions, as depicted in Fig. 3. The distribution of the labels is shown in Table 1.

The WSI images were acquired in 2013–2015 from patients from the Castelo Branco region (Portugal) with a Leica SCN400 (from Centro Hospitalar Cova da Beira), and were made available on .svs format, with a pixel scale of 0.467 $\mu\text{m}/\text{pixel}$ and variable size with width $\in [39980 \ 62952]$ and height $\in [27972 \ 44889]$ (pixels). The ground-truth was released as the coordinates of the points that enclose each labeled region via a .xml file.

2.3. Performance evaluation

The methods developed by the participants were evaluated on independent test sets for which the ground-truth was hidden. Specifically, for Part A participants were requested to submit a .csv containing row-wise pairs of (image name, predicted label) for the 100 microscopy images. Performance on the microscopy images was evaluated based on the overall prediction accuracy, i.e., the ratio between correct samples and the total number of evaluated images.

$\begin{bmatrix} 3 & 2 & 0 & 1 & 2 \\ 3 & 2 & 0 & 1 & 2 \\ 3 & 1 & 0 & 1 & 2 \\ 2 & 1 & 0 & 2 & 3 \end{bmatrix}$	GT	$\begin{bmatrix} 2 & 3 & 0 & 0 & 1 \\ 1 & 2 & 0 & 3 & 0 \\ 0 & 0 & 2 & 3 & 0 \\ 0 & 0 & 3 & 3 & 0 \end{bmatrix}$	$s = 0.56$
$s = 1.0$			$s = 0.33$
$s = 0.89$			$s = 0.08$
$s = 0.56$			$s = 0$

Fig. 4. Examples of the custom score metric. □ | 0 normal; ■ | 1 benign; ■ | 2 *in situ*; ■ | 3 invasive.

For Part B it was required the submission of 4 \times downsampled WSI .png masks with values 0 – Normal, 1 – Benign, 2 – *In situ* carcinoma and 3 – Invasive carcinoma. Possible mismatches between the prediction's and ground truth's sizes were corrected by padding or cropping the prediction masks. The performance on the WSI images was evaluated based on the custom score s :

$$s = 1 - \frac{\sum_{i=1}^N |\text{pred}_i - \text{gt}_i|}{\sum_{i=1}^N \max(\text{gt}_i, |\text{gt}_i - 3|) \times [1 - (1 - \text{pred}_{i,\text{bin}})(1 - \text{gt}_{i,\text{bin}})] + a} \quad (1)$$

where pred is the predicted class (0, 1, 2 or 3), and gt is the ground truth class, i is the linear index of a pixel in the image, N is the total number of pixels in the image, X_{bin} is the binarized value of X , i.e., is 0 if the label is 0 and 1 otherwise, and a is a very small number that avoids division by zero. Fig. 4 illustrates the results of this metric on a set of predictions which are gradually farther from the ground truth.

This score is based on the accuracy metric, aiming at penalizing more the predictions that are farther from the ground truth value. The reasoning behind this metric is the following: in the numerator, the absolute distance between the predicted class and the ground truth is measured for all samples i , which is indicative of how far the prediction is from the ground truth, e.g., if $\text{gt} = 1$ and $\text{pred} = 3$, the distance is 2, whereas if $\text{pred} = 2$ the distance is 1. To normalize these distances, in the first factor of the denominator we consider, for each sample, the largest distance possible having in account the true label, e.g., if $\text{gt} = 0$ the maximum distance possible is 3 ($\max(0, |0 - 3|) = 3$), while if $\text{gt} = 2$ the maximum distance is 2 ($\max(2, |2 - 3|) = 2$). Also, the cases in which the prediction and ground truth are both 0 (Normal class) are not counted, since these can be seen as true negative cases. This is allowed by the second factor of the multiplication of the denominator, which is equal to zero if $\text{pred} = 0$ and $\text{gt} = 0$, and equal to 1 otherwise.

Note that this custom evaluation score was preferred over both the Intersection over Union (IoU) and the quadratic weighted Cohen's Kappa statistic (Cohen, 1960). Namely, the custom score allows to ignore correct Normal class predictions (highly dominant) while penalizing wrong Normal predictions, whereas for Kappa the Normal class would have either to be completely considered or ig-

nored. Likewise, the custom metric is not only able to assess if the methods were properly capable of detecting pathological regions on whole-slide images (as would the IoU, if applied class-wise), but also of indicating how far that prediction is from the ground-truth. Given the complexity of the task, the direct computation of the IoU could over-penalize methods that are capable of finding abnormal regions but fail to correctly classify them. By assessing the distance of the predictions, the custom metric has a higher clinical relevance than analyzing region overlaps. For instance, mispredicting Normal tissue as Benign should be considered less severe than predicting it as Invasive.

3. Competing solutions

This section provides a comprehensive description of the participating approaches. Tables 2 and 3 summarize the methods that achieved an accuracy ≥ 0.7 and score ≥ 0.5 on Part A and B, respectively. The most relevant methods in terms of performance and applicability are detailed on the next sections. For methods that solve Part A and B jointly, refer to Section 3.4, and for Part A or Part B exclusively refer to Sections 3.2 and 3.3, respectively.

3.1. Introduction to convolutional neural networks

The vast majority of Part A and all of Part B participants proposed a convolutional neural network (CNN) approach to solve BACH. CNNs are now the state-of-the-art approach for computer vision problems and show high promise in the field of medical image analysis (Litjens et al., 2017; Tajbakhsh et al., 2016) because they are easy to set up, require little applied field knowledge (specially when compared with handcrafted feature approaches) and allow to migrate base features from generic natural image applications (Deng et al., 2009).

CNN performance is highly dependent on the architecture of the network as well as on the hyper-parameter optimization (learning rate, for instance). The large number of parameters in CNNs make them prone to overfit to the training data, specially when a relatively low number of training images is available. Because of that, it is a common practice in medical image analysis to fine-tune networks trained on medical images. In BACH, participants opted mainly for pre-trained networks that have historically achieved high performance in the ImageNet natural image analysis challenge (Russakovsky et al., 2015). From those, VGG (Simonyan and Zisserman, 2014), Inception (Szegedy et al., 2015), ResNet (He et al., 2016) and DenseNet (Huang et al., 2017) were the ones that achieved the overall higher results. A brief description of these networks is provided below.

VGG (Visual Geometry Group) was one of the first networks to show that increasing model depth allows higher prediction performance. This network is composed of blocks of 2–3 convolutional layers with a large number of filters that are followed by a max pooling layer. The output of the last layer is then connected to a set of fully-connected layers to produce the final classification. However, despite the success of this model on the ImageNet challenge, the linear structure of VGG and large number of parameters (approximately 140M for 16 layers) does not allow to significantly increase the depth of the model and increases tendency to overfit.

The Inception network follows the theory that most activations in a (deep-)CNN are either unnecessary or redundant and thus the number of parameters can be reduced by using locally sparse building blocks (*a.k.a.* inception blocks). At each inception block, the number feature maps of the previous block is reduced via an 1×1 convolution. The projected features are then convolved in parallel by kernels of increasing size, allowing to combine information at multiple scales. Finally, replacing the fully-connected lay-

ers by a global average pooling allows to significantly reduce the model parameters (23M parameters with 159 layers) and makes the network fully convolutional, enabling its application to different input sizes.

Increasing network depth leads to vanishing gradient problems as a result of the large number of multiplication operations. Consequently, the error gradient will be vanishingly small preventing effective updates of the weights in the initial layers of the model. The recent versions of Inception tackle this issue by using Batch Normalization (Ioffe and Szegedy, 2015), which allows to reestablish the gradient by normalizing the intermediary activation maps with the statistics of the training batch. Alternatively, ResNet uses residual blocks to stabilize the value of the error gradient during training. In each residual block the input activation map is summed to the output of a set of convolutional layers, thus stopping the gradient from vanishing and easing the flow of information. A ResNet with 50 residual blocks (169 layers) has approximately 25M parameters.

The high performance of models like Inception or ResNet have strengthened the deep learning design principle that "deeper networks are better" by improving on the feature redundancy and gradient vanishing problems. Recently, an even deeper network, DenseNet, has addressed these same issues by using dense blocks. Dense blocks introduce short connections between convolutional layers, *i.e.*, for each layer, the activations of all preceding layers are used as inputs. By doing so, DenseNet promotes feature re-use, reducing the feature redundancy and thus allowing to decrease the number of feature maps per layer. Specifically, a DenseNet with 201 layers has approximately 20M parameters to optimize.

As already mentioned, fine-tuning of high performance networks trained in natural images is the preferred approach for medical image analysis. Fine-tuning for classification is usually performed as follows: 1) the network is initialized with weights trained to solve a natural image classification problem such as the ImageNet classification task; 2) the classification head, usually a fully-connected layer, is replaced by a new one with randomly initialized parameters; 3) initially, the new classification head is trained for a fixed number of iterations by inputting the medical images and inhibiting the filters of the pre-trained model to change; 4) then, different blocks of the pre-trained model are progressively allowed to learn and adapt to the new features, allowing the model to move to new local optima and increase the overall performance of the network.

3.2. Part A

3.2.1. Chennamsetty et al. (team 216)

(Chennamsetty et al., 2018) used an ensemble of ImageNet pre-trained CNNs to classify the images from Part A. Specifically, the algorithm is composed of a ResNet-101 (He et al., 2016) and two DenseNet-161 (Huang et al., 2017) networks fine-tuned with images from varying data normalization schemes. Initializing the model with pre-trained weights alleviates the problem of training the networks with limited amount of high quality labeled data. First, the images were resized to 224×224 pixels via bilinear interpolation and normalized to zero mean and unit standard deviation according to statistics derived either from ImageNet or Part A datasets, as detailed below.

During training, the ResNet-101 and a DenseNet-161 were fine-tuned with images normalized from the breast histology data whereas the other DenseNet-161 was fine-tuned with the ImageNet normalization. Then, for inference, each model in the ensemble predicts the cancer grade in the input image and a majority voting scheme is posteriorly used for assigning the class associated with the input.

Table 2

Summary of the methods submitted for Part A. ^Adetailed description in Section 3.2; ^{AB}detailed description in Section 3.4. Pre-training is performed on ImageNet (Krizhevsky et al., 2012). **Acc.** is the overall prediction success for the four classes; **Approach** lists the main methods to label the images; **Ensemble** (Ens.) indicates if the approach uses a single or multiple models (and their number, when available); **External sets** indicates if the method was trained using datasets other than from Part A; **Context (area ratio)** is the ratio between the original size and the size of the patch used for training the network (prior to rescaling); **Input size (pixels)** is the size of the image to be analyzed by the model; **Color normalization** (color norm.) indicates if any histology-inspired normalization was used. N/A: information not available. ¹Pre-trained on CAMELYON (<https://camelyon17.grand-challenge.org/>). M1: (Bejnordi et al., 2016); M2: (Krishnan and Shah, 2012); M3: (Macenko et al., 2009); M4: (Reinhard et al., 2001).

Team	Acc.	Approach	Pre-trained	Ens.	External sets	Context (area ratio)	Input size (pixels)	Color norm.
(Chennamsetty et al., 2018) (216) ^A	0.87	Resnet-101; Densenet-161	✓	3	✗	1	224 × 224	✗
(Kwok, 2018) (248) ^{AB}	0.87	Inception-Resnet-v2	✓	✗	Part B	0.71	299 × 299	✗
(Brancati et al., 2018) (1) ^A	0.86	Resnet-34, 50, 101	✓	3	✗	1	308 × 308 615 × 615	✗
(Marami et al., 2018) (16) ^{AB}	0.84	Inception-v3	✓	4	Part B BreakHis	0.33	512 × 512	M1
(Kohl et al., 2018) (54) ^{AB}	0.83	Densenet-161	✓ ¹	✗	✗	1	205 × 154	✗
(Wang et al., 2018a) (157) ^A	0.83	VGG16	✓	✗	✗	0.765	224 × 224	✗
Steinfeldt et al. (186)	0.81	Xception	✓	✗	✗	0.028–0.751	229 × 229	✗
(Koné and Boulmane, 2018) (19) ^A	0.81	ResNeXt50	✓	✓	BISQUE	1	299 × 299	✗
Nedjar et al. (36)	0.81	Inception-v3, Resnet-50, MobileNet	✓	✓	✗	1	224 × 224	✗
Ravi et al. (412)	0.8	Resnet-152	✓	✗	✗	0.875	224 × 224	M2
(Wang et al., 2018b) (22)	0.79	VGG16	✓	✗	✗	0.255	224 × 224	M3
(Cao et al., 2018) (425)	0.79	PTFAS+GLCM, ResNet-18, ResNeXt, NASNet-A, ResNet-152, VGG16, Random Forest SVM	✓	✓	✗	1	224 × 224 331 × 331	✗
Seo et al. (60)	0.79	ResNet, Inception-V3, Random Forests	✓	✗	✓	1	299 × 299	✗
Sidhom et al. (370)	0.78	ResNet-50	✓	✗	✗	0.018 0.289	224 × 224	M4
(Guo et al., 2018) (242)	0.77	GoogLeNet	✓	2	✗	1 0.083	224 × 224	M4
Ranjan et al. (61)	0.77	AlexNet	✓	2	✗	1	224 × 224	✗
(Mahbod et al., 2018) (73)	0.77	ResNet-50, ResNet-101	✓	2	✗	1	224 × 224	M3
(Ferreira et al., 2018) (18)	0.76	Inception-ResNet-v2	✓	✗	✗	1	224 × 224	✗
(Pimkin et al., 2018) (256)	0.76	ResNet34, Densenet169, Densenet201 XGBoost	✓	12	Part B BreakHis	1	300 × 300	✗
Sarker et al. (358)	0.75	Inception-v4	✓	✗	✗	0.083	299 × 299	✗
(Rakhlil et al., 2018) (98)	0.74	VGG16, ResNet-50, IceptionV3, LightGBM	✓	✓	✗	0.20 0.54	400 × 400 600 × 600	M3
(Iesmantas and Alzbutas, 2018) (164)	0.72	Custom CNN (Capsule Network)	✗	✗	✗	0.029	512 × 512	M4
Xie et al. (253)	0.72	CNN	✗	✗	✗	0.083	512 × 512	✗
(Weiss et al., 2018) (268)	0.72	Xception, Logistic Regression	✓	✗	✗	1	1024 × 768	M3
(Awan et al., 2018) (6)	0.71	ResNet50, SVM	✓	✗	✗	0.33	512 × 512	M4
Liang (62)	0.7	VGG16, VGG19, ResNet50, InceptionV3, Inception-Resnet, k-NN	✓	5	✗	0.083	N/A	✗

Table 3
 Summary of the methods submitted for whole-slide image analysis (Part B). ^aDetailed description in Section 3.3; ^bdetailed description in Section 3.4. Pre-training is performed on ImageNet (Krizhevsky et al., 2012) unless stated otherwise. **Score** is the custom metric from Eq. (1); **Approach** lists the main methods to label the images; **Ensemble** (Ens.) indicates if the approach uses a single or multiple models (and their number, when available); **External sets** indicates if the method was trained using datasets other than from Part A; **Context (area ratio)** is the ratio between the average original size and the size of the patch that is used for training the network (prior to rescaling); **Input size (pixels)** is the size of the image to be analyzed by the model; **Color normalization** (Color norm.) indicates if any histology-inspired normalization was used. ¹ trained on ImageNet, ² trained on VOC2012.

Team	Score	Approach	Pretrained	Ens.	External sets	Context (area ratio)	Input size (pixels)	Color norm.
(Kwok, 2018) (248) ^{aB} (Marani et al., 2018) (16) ^{aB}	0.69 0.55	Inception-Resnet-v2 Inception-v3 + adaptive pooling	✓ ✓	× 4	Part A Part A BreakHis	8.5e-4 9.9e-5	299 × 299 512 × 512	× ×
Jia et al. (296)	0.52	ResNet-50 + multiscale atrous convolution	✓	×	×	9.9e-5	512 × 512	×
Li et al. (137)	0.52	VGG16 ¹ , DeeplabV2 ¹ , Resnet50 ²	✓	✓	×	9.9e-5	512 × 512	×
Murata et al. (91)	0.50	U-Net	×	×	×	1.6e-2	256 × 256	×
(Galal and Sanchez-Freire, 2018) (264) ^B	0.50	DenseNet	×	×	×	1.6e-3	2048 × 2048	×
(Vu et al., 2018) (166) ^{aB}	0.49	DenseNet, SENet, ResNext	✓	×	×	1.5e-4	630 × 630	×
(Kohl et al., 2018) (54) ^{aB}	0.42	Densenet-161	✓	×	Part B non-annotated	9.3e-6	157 × 157	×

3.2.2. Brancati et al. (team 1)

Brancati et al. (2018) proposed a deep learning approach based on a fine-tuning strategy by exploiting transfer learning on an ensemble of ResNet (He et al., 2016) models. ResNet was preferred to other deep network architectures because it has a small number of parameters and shows a relatively low complexity in comparison to other models. The authors opted by further reducing the complexity of the problem by down-sampling the image by factor k and using only the central patch of size $m \times m$ as input to the network. In particular, k was fixed to 80% of the original image size and m was set equal to the minimum size between the width and high of the resized image.

The proposed ensemble is composed of 3 ResNet configurations: 34, 50 and 101. Each configuration was trained on the images from Part A and the classification of a test image is obtained by computing the highest class probability provided by the three configurations.

3.2.3. Wang et al. (team 157)

Wang et al. (2018a) proposed the direct application of VGG-16 (Simonyan and Zisserman, 2014) to solve Part A. Prior to fine-tuning the model, all images from Part A were resized to 256×256 and normalized to zero mean and unit standard deviation. To account for the model input size, training is performed by cropping patches of 224×224 pixels at random locations of the input image. First, the model is trained using a Sample Pairing (Inoue, 2018) data augmentation scheme. Specifically, a random pair of images of different labels is independently augmented (translations, rotations, etc.) and then superimposed with each other. The resulting mixed patch receives the label of one of the initial images and is afterwards used to train the classifier. The learned weights are then used as a starting point to train the network with the initial (*i.e.* non mixed) dataset.

3.2.4. Kone et al. (team 19)

Koné and Boulmane (2018) proposed a hierarchy of 3 ResNeXt50 (Xie et al., 2017) models in a binary tree like structure (one parent and two child nodes) for the 4-class classification of Part A. The top CNN classifies images in two high level groups: 1. carcinoma, which includes the *in situ* and the invasive classes and 2. non-carcinoma, which includes normal and benign. Then, each of children CNNs sub-classifies the images in the respective 2 classes.

The training is performed in two steps. First, the parent ResNeXt50 pre-trained on ImageNet is fine tuned with the images from Part A. The learned filters are then used as the starting point for the child networks. The authors also divide the ResNeXt50 layers into three groups and assign them different learning rates based on the optimal one found during training.

3.3. Part B

3.3.1. Galal et al. (team 264)

Galal and Sanchez-Freire (2018) proposed Candy Cane, a fully convolutional network based on DenseNets (Huang et al., 2017) for the segmentation of WSIs. Candy Cane was designed following an auto-encoder scheme of downsampling and upsampling paths with skip connections between corresponding down and up feature maps to preserve low level feature information. Accounting for GPU memory restrictions, the authors propose a downsampling path much longer than the upsampling counter-part. Specifically, Candy Cane operates on 2048×2048 slice images and outputs the corresponding labels at a size of 256×256 pixels. Similarly to an expert that looks at a microscope in few adjacent regions to examine the tissue but then identifies regions in the larger context of the tissue, the large input size of the model allows the network

to have both microscopy and tissue organization contexts. The output of the system is then resized to the original size.

3.4. Part A and B

3.4.1. Kwok et al. (team 248)

Kwok (2018) used a two-stage approach to take advantage of both microscopy and WSI images. To account for the partially missing patient-wise origin on Part A, images which origin was not available were clustered based on color similarity. The data was then split accordingly. For stage 1, 5600 patches of 1495×1495 pixels were extracted from Part A's images with a stride of 99 pixels. These patches were then resized to 299×299 pixels and used for fine-tuning a 4 class Inception-Resnet-v2 (Szegedy et al., 2017) trained on ImageNet. This network was then used for analyzing the WSIs. Specifically, WSI foreground masks were computed via a threshold on the L^*a^*b color space. Then, patches were extracted from the WSIs in the same way as for Part A. This second patch dataset was refined by discarding images with $< 5\%$ foreground and posteriorly labeled using the CNN trained on Part A. Finally, 5900 patches from the top 40% incorrect predictions (evenly sampled from each of the 4 classes) were selected as hard examples for stage 2.

For stage 2, the CNN was retrained by combing the patches extracted from Part A (5600) and Part B (5900). The resulting model was used for labeling both microscopy images and WSIs. Prediction results were aggregated from patch-wise predictions back onto image-wise predictions (for Part A) and WSI-wise heatmaps (for Part B). Specifically for Part B, the patch-wise predictions were mapped to hard labels (Normal=0, Benign=1, *in situ*=2 and Invasive=3) and combined into a single image based on the patch coordinates and the network's stride. The resulting map was then normalized to [0 1] and multi-thresholded at {0, 0.35, 0.7, 0.75} to bias the predictions more towards Normal/Benign and less to *in situ* and Invasive carcinomas.

3.4.2. Marami et al. (team 16)

Marami et al. (2018) proposed a classification scheme based on an ensemble of four modified Inception-v3 (Szegedy et al., 2016) CNNs that aims at increasing the generalization capability of the method by combining different networks trained on random subsets of the data. Specifically, the networks were adapted by adding an adaptive pooling before a set of custom fully connected layers, allowing higher robustness to small scale changes. Each of these CNNs was trained via a 4 fold cross-validation approach on 512×512 images extracted at $20 \times$ magnification from both microscopy images from Part A and WSI from Part B, as well as with benign tissue images from the BreakHis public dataset (Spanhol et al., 2016b).

Predictions on unseen data are inferred by averaging the output probabilities of the trained ensemble network for each class, making the system more robust to potential inconsistencies and corruption in the labeled data. For Part A, the final label was obtained by majority voting of 12 overlapping 512×512 regions. For the WSIs, local predictions were generated by using a 512×512 sliding window of stride 256 pixels. The resulting output map was then refined using a ResNet34 (He et al., 2016) to separate tissue regions from the background and regions with artifacts, reducing potential misclassifications due to ink and other artifacts in whole slide images.

3.4.3. Kohl et al. (team 54)

Kohl et al. (2018) used an ImageNet pretrained DenseNet (Huang et al., 2017) to approach both parts of the challenge. For Part A, the 400 training images were downsampled by a factor 10 and normalized to zero mean and unit standard deviation. The

network was then trained on two steps: 1) fine-tuning the fully-connected portion of the network by 25 epochs to avoid over-fitting and 2) training the entire network for 250 epochs.

For Part B, the authors extracted patches of $330 \mu\text{s} \times 330 \mu\text{s}$ (157×157 pixels) from the annotated WSIs. This patch dataset was then refined by removing patches consisting of at least 80% background pixels, similarly to Litjens et al. (2016). Due to the very limited amount of data in the benign and *in situ* carcinoma classes, the authors did not perform WSI-wise split for validation purposes and instead used a randomly split dataset. Also, 16 of the 20 originally non-annotated WSIs were also annotated with the help of a trained pathologist and thus in total 41,506 image patches (25,230 normal, 1723 benign, 1759 *in situ* and 12,794 invasive carcinomas) were used. Network training was similar for Part A and Part B: 25 epochs for training the fully-connected layers followed by 250 epochs for training the whole network in case of Part A, and 6 epochs for training the fully-connected layers followed by 100 epochs for training the whole network using log-balanced class weights in case of Part B.

3.4.4. Vu et al. (team 166)

Vu et al. (2018) proposed to use an encoder-decoder network to solve both Part A and Part B. For Part A, the authors use the encoder part of the model. The encoder is composed of five convolutional processing blocks that integrate dense skip connections, group and dilated convolutions, and self-attention mechanism for dynamic channel selection following the design trends of DenseNet, Squeeze-Excitation Network (SENet) and ResNext (Huang et al., 2017; Jégou et al., 2016; Yu and Koltun, 2015; Hu et al., 2017; Xie et al., 2017). For classifying the microscopy images, the model has a head composed of a global average pooling and a fully-connected softmax layer. Training is performed by downsampling the images $4 \times$ and online data augmentation is used.

For Part B the full encoder-decoder scheme is used. This segmentation network follows the U-Net (Ronneberger et al., 2015) structure with skip connections between the downsample and up-sample, the decoder is composed by the same convolutional blocks and the upsample is performed considering the nearest neighbor. Also, to ease network convergence, the encoder is initialized with the weights learned from Part A. For training, the WSI are first downsampled by a factor of 4 and sub-regions containing the labels of interest are collected. Specifically, the authors collect 6129 sub-regions of size 1000×1000 from which the central regions of 630×630 are used as input to the model. The corresponding output segmentation map has a size of 204×204 . To prioritize the detection of pathological regions, the segmentation network is trained with two categorical crossentropy loss terms, where the main loss targets for the four histology classes and the auxiliary loss is computed for normal and benign vs *in situ* carcinoma and invasive carcinoma groups.

4. Results

The BACH had worldwide participation, with a total of 677 registrations and 64 submissions for both part A (51) and B (13), as shown on Fig. 5.

4.1. Performance in Part A

Participants of the Part A of the challenge were ranked in terms of accuracy. As in Araújo et al. (2017), these submissions were further evaluated in terms of sensitivity and specificity:

$$\text{sensitivity} = \frac{TP}{TP + FN} \quad (2)$$

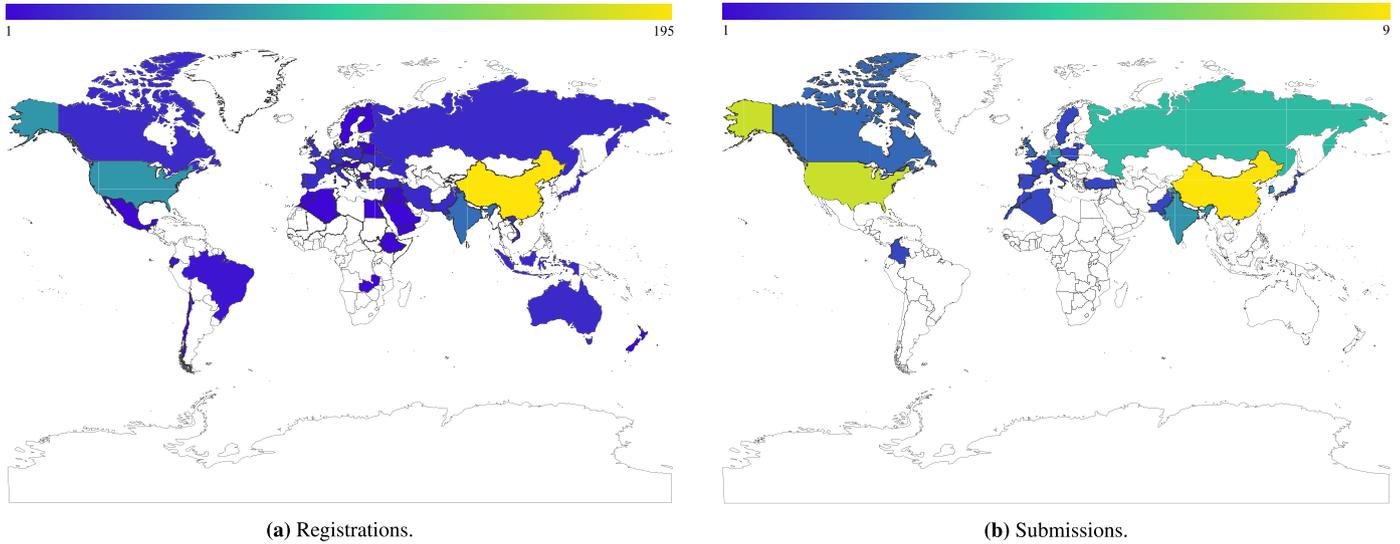


Fig. 5. Geographical distribution of the BACH participants. a) registered on the website; b) submitted for the test set.

$$\text{specificity} = \frac{TN}{TN + FP} \quad (3)$$

where TP , TN , FP and FN are the class-wise true-positive, true-negative, false-positive and false-negative predictions, respectively. For benchmarking purposes, a simple fine-tuning experiment on the BACH Part A was conducted. Specifically, the classification parts of VGG16, Inception v3, ResNet50 and DenseNet169 were replaced by a pair of untrained fully-connected layers with 1024 and 4 neurons. These networks were then trained on two steps, first by updating only the new fully-connected layers until the validation loss stopped improving and posteriorly training the entire model until the same stop criteria was met. Adam (Kingma and Ba, 2015) was used as optimizer and the loss was the categorical cross-entropy.

Finally, to further evaluate the performance of the methods submitted to Part A, four pathologists (E1–E4) were asked to classify the field images from the BACH test set. E1 and E3 are breast cancer specialists and E2 and E4 are experienced pathologists. Also, E1 was one of the two experts involved in the construction of training and testing sets from BACH, and the remaining three were external to the process. The difference in the annotation process was that in the BACH sets construction the pathologists had access to other regions of the patient tissue (and potentially immunohistochemical analysis), whether in this second phase they could only see the field image, *i.e.*, they only had access to the same information of the automated classification algorithms.

The class-wise performance of the methods is shown in Table 4. Table 5 shows the performance for two binary cases: 1) a referral scenario, *Pathological*, where the objective is to distinguish Normal images from the remaining classes and, 2) a cancer detection scenario, *Cancer*, where the Normal and Benign classes are grouped vs the *In situ* and Invasive classes.

Fig. 6a depicts, for the top-10 participants, the difference between the reported performances on the training set (cross-validation) and the achieved performances on the hidden test set. Also, a class-wise study of these methods shows that the Benign and *In situ* classes are the most challenging to classify (Fig. 6b). In particular, Fig. 6c and d shows two images with 100% inter-observer agreement that were misclassified by the majority (at least 80%) of the top-10 approaches.

4.1.1. Inter-observer analysis

The accuracies of the three external pathologists are of 94%, 78% and 73%, and the accuracy of the pathologist from BACH is 96%. Note that this pathologist annotated the images after one month from the first annotation, in order to avoid the influence of past knowledge regarding the patients' exams. For comparison purposes, Fig. 7 shows the inter-observer, inter-method and observer-method agreement via the quadratic-weighted Cohen's kappa score and their corresponding statistical differences, and Fig. 8 shows the confusion matrices of the experts.

4.2. Performance in Part B

The overall challenge performance and main approaches of the participating teams are shown in Table 3. Similarly to Part A, please refer to Tables 6 and 7 for the team-wise sensitivity and specificity of the methods. For reference purposes, Table 6 also shows the quadratic-weighted kappa scores for each method. Examples of pixel-wise predictions for Part B are shown in Fig. 9. The correct identification of invasive regions was more successful, as opposed to benign and *in situ* regions.

4.3. Statistical analysis

The top-10 methods of Part A and the annotations of the medical experts are statistically compared via an adaption of McNemar's test (Edwards, 1948; Dietterich, 1993). The McNemar test is based on the chi-squared distribution and allows to assess the performance of two classifiers based on their accuracy on an independent test set. Let A and B be two methods to compare. The chi-squared (χ^2) distribution with 1 degree of freedom is defined by Eq. (4):

$$\chi^2 = \frac{(|n_{01} - n_{10}| - 1)^2}{n_{01} + n_{10}} \quad (4)$$

where n_{01} is the number of misclassified test samples by B but not by A and n_{10} is the number of samples misclassified by A but not by B. The null assumption that A and B have equal classification performance is rejected if $\chi^2 > 3.841$, corresponding to a p -value of 0.05. The statistical analysis for Part A is summarized in Fig. 7.

The Part B's submissions are statistically assessed by the confidence intervals of their average performance in terms of the BACH score s and quadratic-weighted kappa score. Namely, assuming

Table 4

Class-wise sensitivity and specificity of Part **A** approaches for the classes in study. Benchmarking results via fine-tuning are also shown. Acc - accuracy; Se - sensitivity; Sp - specificity. Expert 1 annotated the BACH dataset.

Team	Acc	Normal		Benign		<i>In situ</i>		Invasive	
		Se.	Sp.	Se.	Sp.	Se.	Sp.	Se.	Sp.
216	0.87	0.96	0.88	0.8	0.96	0.84	1.0	0.88	0.99
248	0.87	0.96	0.93	0.72	0.96	0.88	0.97	0.92	0.96
1	0.86	0.96	0.91	0.68	0.97	0.84	0.99	0.96	0.95
16	0.84	0.92	0.95	0.64	0.96	0.84	0.99	0.96	0.89
54	0.83	0.96	0.92	0.52	0.97	0.88	0.92	0.96	0.96
157	0.83	0.96	0.91	0.64	0.99	0.92	0.91	0.8	0.97
186	0.81	0.96	0.92	0.68	0.96	0.76	0.95	0.84	0.92
19	0.81	1.0	0.95	0.4	0.99	0.92	0.92	0.92	0.89
36	0.81	0.88	0.92	0.6	0.96	0.88	0.95	0.88	0.92
412	0.8	0.92	0.96	0.48	0.97	0.84	0.92	0.96	0.88
VGG	0.58	0.84	0.84	0.64	0.84	0.72	0.87	0.36	0.97
Inception	0.77	0.92	0.93	0.44	0.96	0.88	0.87	0.84	0.93
ResNet	0.76	0.88	0.92	0.52	0.95	0.8	0.87	0.84	0.95
DenseNet	0.79	0.92	0.96	0.36	0.99	0.92	0.83	0.96	0.95
Expert 1	0.96	0.96	0.99	0.92	0.97	1.0	1.0	0.96	0.99
Expert 2	0.94	0.96	0.99	0.88	0.96	1.0	1.0	0.92	0.97
Expert 3	0.78	0.88	0.99	0.76	0.79	0.56	0.97	0.92	0.96
Expert 4	0.73	0.40	0.99	0.84	0.71	0.76	0.97	0.92	0.97
Experts (avg.)	0.85 ± 0.10	0.80 ± 0.23	0.99 ± 0.00	0.85 ± 0.06	0.86 ± 0.11	0.83 ± 0.18	0.99 ± 0.01	0.93 ± 0.02	0.97 ± 0.01

Table 5

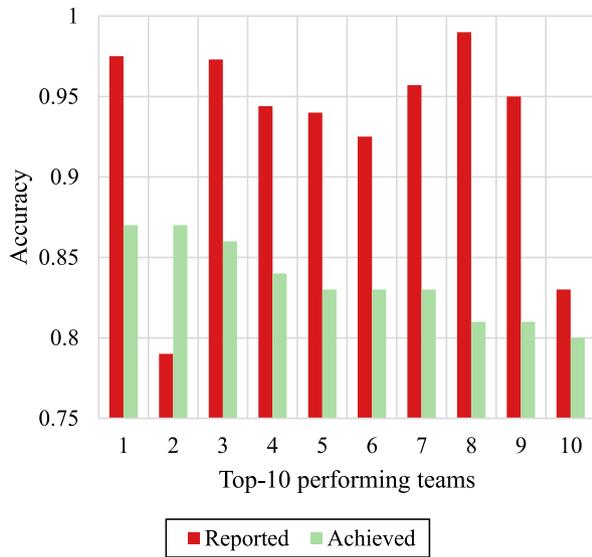
Class-wise sensitivity and specificity of Part **A** approaches. **Patho.** refers to Benign, *in situ* and Invasive vs Normal, and **Cancer** refers to *in situ* and Invasive vs Normal and Benign classes. Benchmarking results via fine-tuning are also shown. Acc - accuracy (\pm the confidence interval); Se - sensitivity; Sp - specificity. Expert 1 annotated the BACH dataset.

Team	Patho.			Cancer		
	Acc.	Se.	Sp.	Acc.	Se.	Sp.
216	0.9	0.88	0.96	0.92	0.86	0.98
248	0.94	0.93	0.96	0.92	0.92	0.92
1	0.92	0.91	0.96	0.9	0.9	0.9
16	0.94	0.95	0.92	0.9	0.94	0.86
54	0.93	0.92	0.96	0.89	0.94	0.84
157	0.92	0.91	0.96	0.94	0.96	0.92
186	0.93	0.92	0.96	0.9	0.9	0.9
19	0.96	0.95	1.0	0.86	0.96	0.76
36	0.91	0.92	0.88	0.86	0.9	0.82
412	0.95	0.96	0.92	0.86	0.96	0.76
VGG16	0.84	0.84	0.84	0.66	0.66	0.88
Inception V3	0.93	0.93	0.92	0.92	0.92	0.76
ResNet 50	0.92	0.92	0.88	0.86	0.86	0.76
DenseNet 169	0.96	0.96	0.92	0.96	0.96	0.68
Expert 1	0.98	0.99	0.96	0.98	0.98	0.98
Expert 2	0.98	0.99	0.96	0.96	0.96	0.96
Expert 3	0.96	0.99	0.88	0.82	0.74	0.90
Expert 4	0.84	0.99	0.40	0.90	0.86	0.94
Experts (avg.)	0.94 ± 0.06	0.99 ± 0.00	0.80 ± 0.24	0.91 ± 0.06	0.88 ± 0.09	0.95 ± 0.03

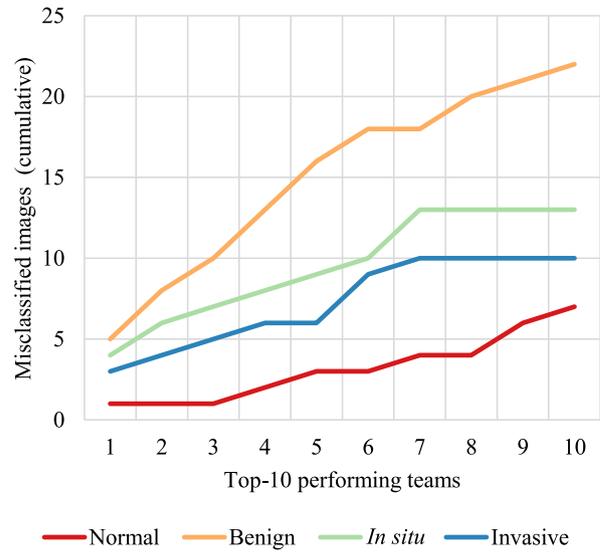
Table 6

Class-wise sensitivity and specificity of Part **B** approaches for the classes in study, along with the challenge score and the quadratic-weighted kappa score (\pm the confidence interval). Se - sensitivity; Sp - specificity; Score - challenge score s (Eq. (1)); k - kappa score.

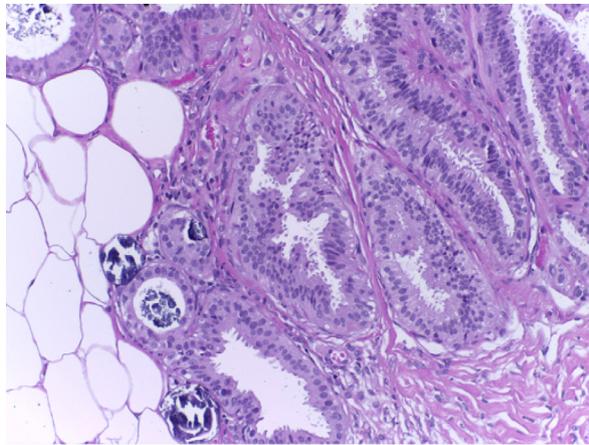
Team	Score	k	Benign		<i>In situ</i>		Invasive	
			Se.	Sp.	Se.	Sp.	Se.	Sp.
248	0.69 ± 0.07	0.44 ± 0.18	0.36	0.7	0.03	0.59	0.4	0.96
16	0.55 ± 0.11	0.51 ± 0.15	0.09	0.99	0.05	0.95	0.45	0.92
296	0.52 ± 0.16	0.48 ± 0.19	0.07	0.99	0.03	0.95	0.39	0.93
137	0.52 ± 0.15	0.51 ± 0.18	0.04	1	0.02	0.92	0.53	0.89
91	0.50 ± 0.10	0.28 ± 0.12	0.05	0.8	0.18	0.53	0.13	0.89
264	0.50 ± 0.13	0.29 ± 0.17	0.05	0.88	0.08	0.52	0.47	0.74
166	0.49 ± 0.11	0.28 ± 0.15	0.14	0.9	0.05	0.63	0.44	0.76
94	0.47 ± 0.15	0.39 ± 0.18	0.16	0.95	0.05	0.76	0.5	0.78
256	0.46 ± 0.14	0.17 ± 0.14	0.18	0.58	0	0	0.58	0.47
15	0.46 ± 0.13	0.27 ± 0.15	0.11	0.78	0.02	0.46	0.4	0.68
54	0.42 ± 0.14	0.34 ± 0.15	0.03	0.98	0.06	0.75	0.52	0.74
183	0.39 ± 0.10	0.12 ± 0.13	0.31	0.81	0.15	0.47	0.15	0.71
252	0.33 ± 0.12	0.13 ± 0.09	0.02	0.98	0	0.85	0.2	0.88



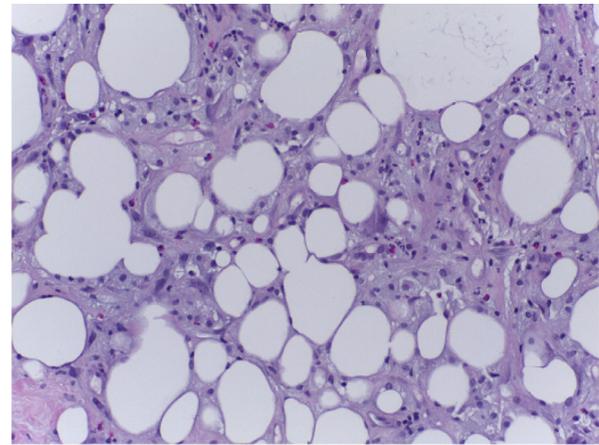
(a) Reported performance during the method submission and respective test accuracy for the top-10 performing teams.



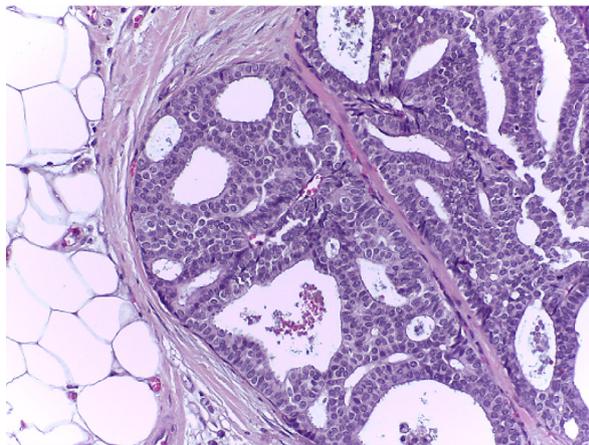
(b) Cumulative number of unique misclassifications per class for the top-10 performing teams. Higher values are indicative of a more challenging classification.



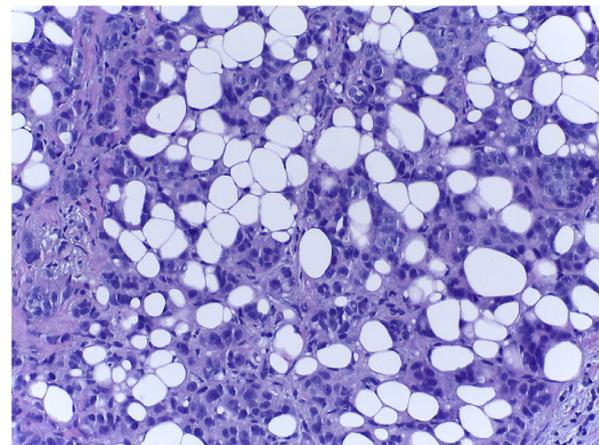
(c) Case of Benign mostly misclassified as *In situ*



(d) Case of Benign mostly misclassified as Invasive



(e) Example of an *In situ* case from the training set



(f) Example of an Invasive case from the training set

Fig. 6. Examples of images misclassified by the top-10 methods of Part A and similar examples in the training set.

that the scores belong to a Gaussian distribution, the confidence interval ci is computed as in Eq. (5):

$$ci_m = \pm 1.96 \frac{\sigma_m}{\sqrt{n}} \quad (5)$$

where 1.96 is the critical value of the confidence interval for $p = 0.05$, σ_m is the standard deviation of one of the studied scores for method m and n is the size of the population (10 for this study).

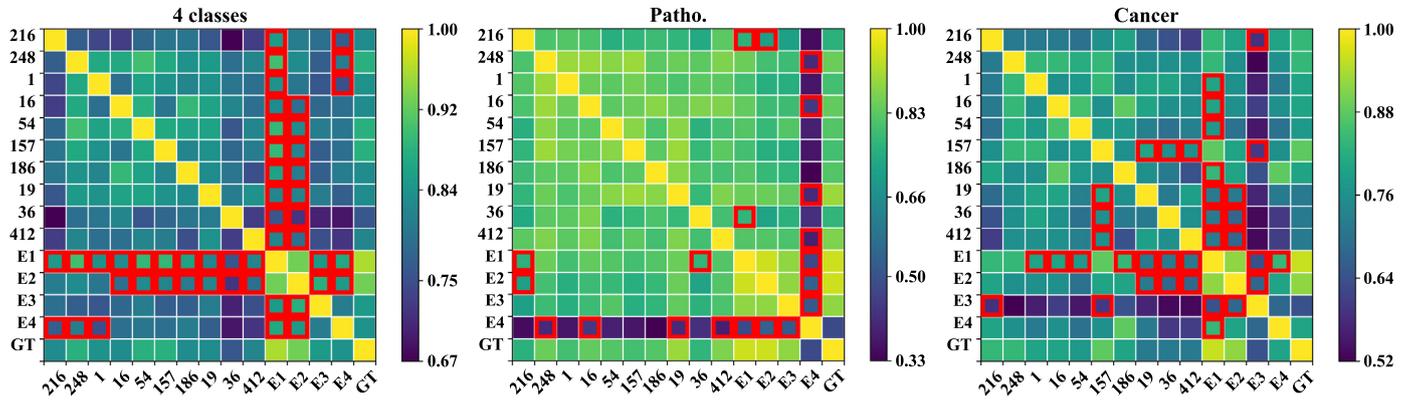


Fig. 7. Inter-observer, inter-method and observer-method quadratic-weighted kappa score Part A. E1–E4 are four expert pathologists, GT is the ground-truth and the numbers indicate the respective team. Expert 1 annotated the BACH dataset one month before participating in the inter-observer study. **4 classes** considers the 4 classes used in this study. **Patho.** refers to Benign, *In situ* and Invasive vs Normal, and **Cancer** refers to *In situ* and Invasive vs Normal and Benign classes. □ the prediction is statistically different according to the McNemar’s test.

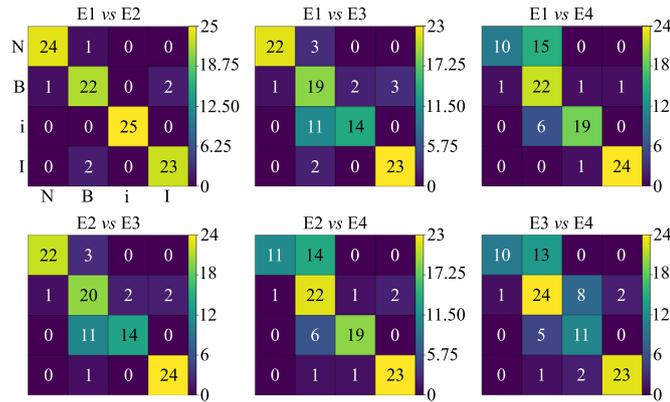


Fig. 8. Confusion matrices of the 4 expert annotators (E#) for the 100 test images of Part A. N - Normal; B - Benign; i - *in situ*; I - Invasive.

Table 7

Class-wise sensitivity and specificity of Part B approaches. **Patho.** refers to Benign, *In situ* and Invasive vs Normal, and **Cancer** refers to *In situ* and Invasive vs Normal and Benign classes. Se - sensitivity; Sp - specificity.

Team	Patho.		Cancer	
	Se.	Sp.	Se.	Sp.
248	0.78	0.59	0.46	0.93
16	0.6	0.95	0.52	0.93
296	0.53	0.95	0.43	0.93
137	0.63	0.92	0.58	0.89
91	0.71	0.53	0.55	0.71
264	0.81	0.52	0.61	0.58
166	0.68	0.63	0.52	0.72
94	0.7	0.76	0.58	0.78
256	0.9	0	0.68	0.48
15	0.82	0.46	0.56	0.65
54	0.68	0.75	0.57	0.74
183	0.8	0.47	0.45	0.61
252	0.3	0.85	0.29	0.87

5. Discussion

BACH accounted for a large number of final submissions in comparison to other medical image challenges. Despite this, there is a similar significant difference between the number of registrations and effective submissions. This stems from common factors such as 1)registrations to inspect the data before deciding to par-

ticipate or to get the data for other purposes, 2) difficulty in downloading the data, which is specially true in countries with Internet accessibility limitations, and 3) high complexity of the task, specially of part B. The verified drop on the submission rate is common on biomedical imaging challenges,⁸ which points to a need to revise future challenge designs to keep the participants’ interest throughout its entire duration. BACH, similarly to other medical image challenges, partially addressed this issue by partnering with the ICIAR conference, which empirically motivated participants by providing an opportunity to show their developed work to the scientific community. For future challenge organizations, it will be needed to further promote participation not only by improving data access but also, for instance, establishing intermediary benchmark timepoints in which participants can compare their performance to motivate themselves to further improve their methods.

The vast majority of the submitted methods used deep learning for solving both tasks A and B. This follows the common trend on the field of medical image analysis, where deep learning approaches are complementing or even replacing the standard manual feature engineering approaches since they allow to achieve high performances while significantly reducing the need for field-knowledge (Litjens et al., 2017). As known, deep learning approaches require large amounts of training data to produce a generalizable model, which are usually not available for medical image analysis due to complexity and high cost of the annotation process. As a consequence, it is a common practice to initialize the models with filters trained on large datasets of natural images, such as ImageNet, and fine-tune them to the specific problem (Tajbakhsh et al., 2016). In fact, as shown in Table 2, all of the top performing methods are composed by one or more deep CNNs architectures such as Inception (Szegedy et al., 2015), DenseNet (Huang et al., 2017), VGG (Simonyan and Zisserman, 2014) or ResNet (He et al., 2016) pre-trained on ImageNet. The difference in performance is thus mainly a consequence of design and training details. For Part A, and unlike previous approaches for the analysis of breast histology cancer images (Araújo et al., 2017), the results of BACH suggest that training the models with a large portion/entire image (even if resized to fit the standard input size of the network) is better than using local patches. This indicates that the overall nuclei and tissue organization may be more relevant than nuclei-scale features, such as nuclei texture, for distinguishing different types of breast cancer. Interestingly this matches the

⁸ <https://grand-challenge.org/challenges>.

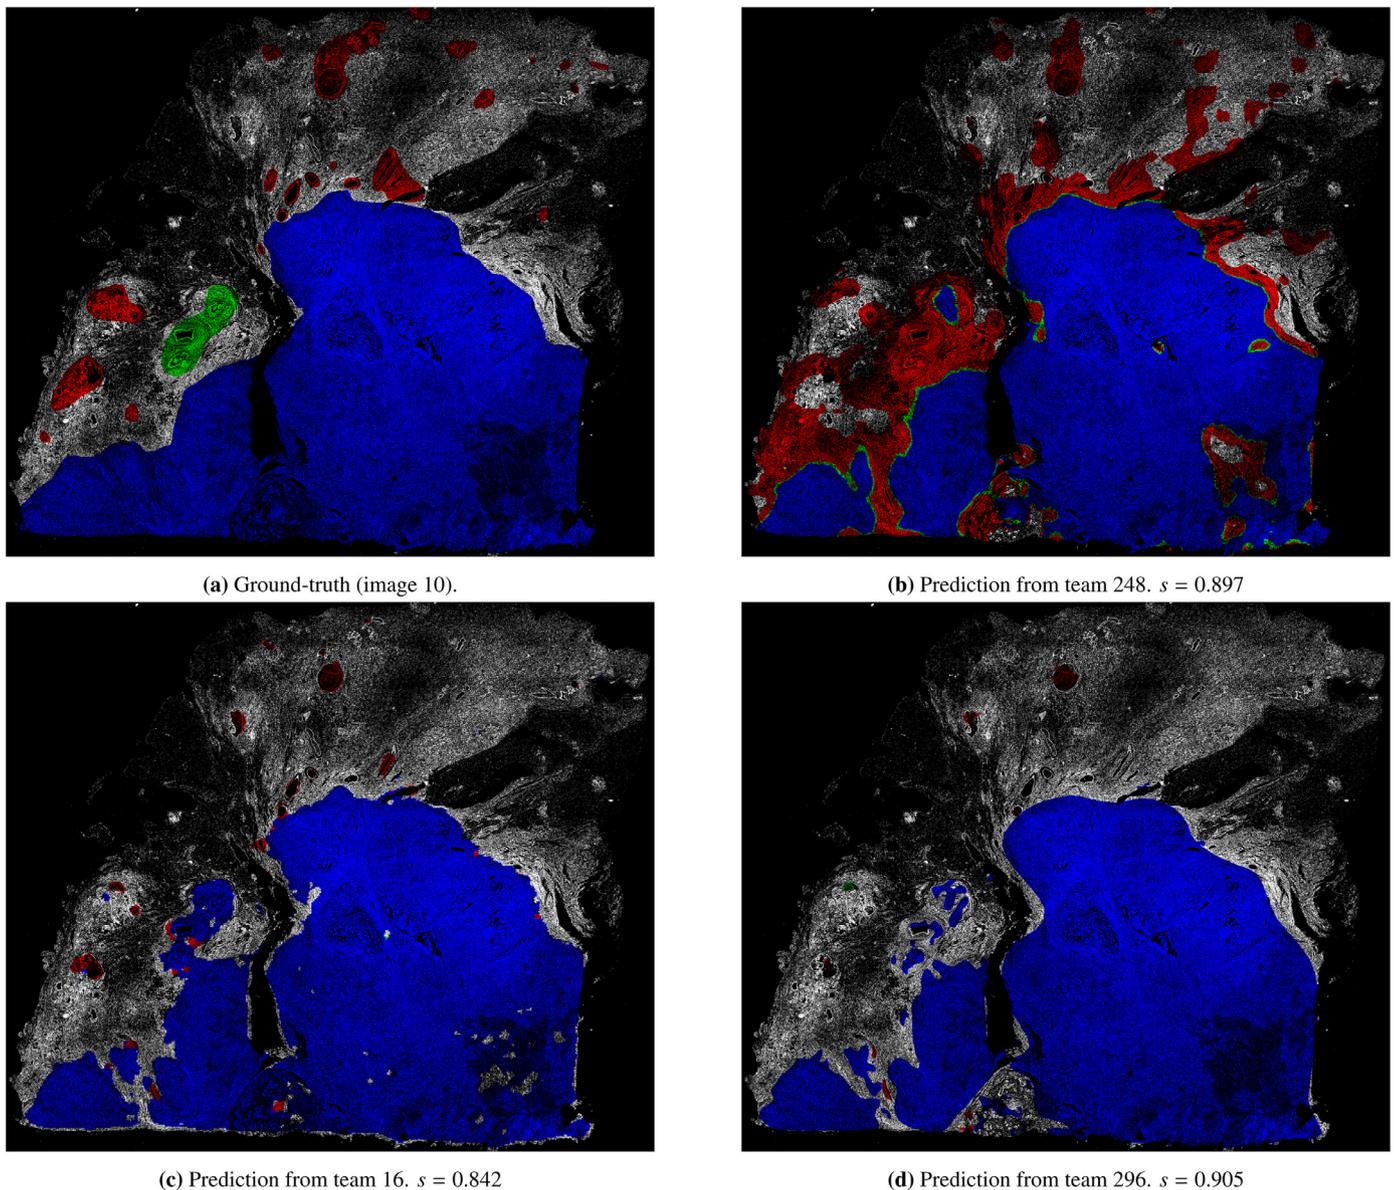


Fig. 9. Examples of test set predictions for Part B from the top performing teams. ■ benign; ■ *in situ*; ■ invasive. The original WSIs were converted to grayscale and the teams' predictions overlaid (background was removed, appearing in black). The obtained s scores (Eq. (1)) are also shown.

importance that clinical pathologists give to tissue architecture features in the diagnostic task. In fact, unlike small patch-based approaches, using large portions of the images eases the integration of both local and global context in the decision process. Besides, patch-based approaches have to handle the problem of patch-label attribution based on the image-level label. Although more sophisticated methods, such as Multiple Instance Learning-based ones could be applied, the vast majority of the teams attributes the label of the image to the patch, which has obvious limitations since the patch may contain only normal tissue, for instance, and be labeled with a different class.

For Part B, the large image size inhibits the direct application of standard segmentation networks, such as U-Net (Ronneberger et al., 2015), to the entire image. Consequently, participants dealt with the issue by analyzing local patches and performing a posterior fusion of the outputs to produce the final probability mapping. In fact, following the same trend of Part A, these methods preferred a large receptive field that guarantees the inte-

gration of contextual and local features during prediction and thus eases the generation of the final class map.

5.1. Performance in Part A

A significant number of submitted methods surpassed the performance of Araújo et al. (2017), which reported an overall 4-class accuracy of 77.8%. Indeed, BACH provided a larger and more representative dataset which, when combined with advances on architectures and transfer-learning techniques, has enabled the development of methods with higher generalization ability. Specifically, these architectures show a high sensitivity for cancer (specially Invasive) cases, which are of great relevance in terms of clinical application (faster automated diagnosis in the cases demanding more urgent attention). Also, as depicted in Tables 4 and 5, the approaches proposed by the participants outperform simple fine-tuning solutions, indicating that there was a clear effort to improve network performance by changing relevant design and training

details. Even though the performance of the top-10 methods is not statistically different (see Fig 7), a careful design of the experimental setting, including train-validation split, data augmentation, architecture combination and parameter tuning, are essential to increase the performance of deep learning systems.

Despite their high accuracy, the submitted methods still failed on correctly predicting images of the more subtle Benign and *In situ* classes. In fact, Fig. 5b shows that the Benign class is the one that affects the most the performance of the methods, which is to be expected since the presence of normal elements and usual preservation of tissue architecture associated with benign lesions makes this class specially hard to distinguish from normal tissue. Furthermore, the Benign class is the one that presents greater morphological variability and thus discriminant features are more difficult to learn.

The generalization capacity of the methods is also affected by the image acquisition pipeline. Specifically, during the acquisition of field images, pathologists focus on capturing regions that contain representative features (tissue architecture, cytological features, etc.) for the given label. As a consequence, whenever those features are subtle, as it is common on normal tissue, specialists tend to capture non-relevant structures, such as fat cells. Likewise, for *in situ* carcinomas it is common to center the images on mammary ducts, where the cancer is contained. Fig. 6c and e shows two images from the test and training sets, respectively. Fig. 6e, labeled as *In situ*, is centered on a duct and surrounded on the left by non-relevant fat tissue. Fig. 6c has, by coincidence, the same acquisition scheme of Fig. 6e and despite being correctly classified as Benign by 100% of the experts, 60% of the top-10 methods classified it as *In situ* and 10% as Normal. Likewise, Fig. 6d shows a full-consensus Benign test image that was classified by 70% of the top-10 as Invasive. Once again, this image has a similar overall tissue organization as training cases of other classes, as shown in the invasive tissue depicted in Fig. 6f. The differences, which lie in the cytological features (nuclei size, color and variability), are clear and yet the approaches failed to correctly capture these discriminant characteristics. This suggests that the networks may be partially modeling *how* images were acquired instead of focusing on *what* leads to the classification (Abràmoff et al., 2016).

Finally, Fig. 6a shows the difference between the top-10 participants' results reported at the submission time via splitting of the training data (e.g. cross-split as train-validation-test) with the achieved performance on the independent test set. The majority of the methods has a 10% difference over the expected accuracy, showing how important a proper design of a method's evaluation design is (and how cross-validation scores can be overly optimistic). Namely, this difference may be due to: 1) patient-wise overfit to the training data, i.e., the authors did not had in account the origin of the images when doing the split and, due to the lack of staining normalization, the networks may have memorized specific staining patterns. In fact, Kwok (2018) was the only top-10 performer to report a lower expected accuracy. As described in Section 3.4, the author performed patient-wise division by clustering images of similar colors, which may contributed to the robustness of the method; 2) over-optimistic train-validation-test split by doing a single split round; and/or 3) excessive hyper parameter-tuning to increase the performance on the split test set, reducing generalization capability. With this in mind, future versions of BACH will provide guidelines on data splitting to reduce this discrepancy and improve the overall scientific correctness of the reported results.

5.1.1. Inter-observer analysis

The Part A BACH dataset was manually annotated by two medical experts and images with dubious diagnosis were discarded. As expected, the annotator of the dataset tends to be better than

his/her peers, which were not capable of correctly classifying, in average, more than 82% of all images. Consequently, in the best scenario the performance of the automatic methods is expected to be equal to that of the observers. Taking into account the average expert accuracy of $85 \pm 10\%$, one can see that the performance obtained by the competing solutions (Table 2) is in line with this value, being the highest accuracy of 87%. The human-level performance of the algorithms is further corroborated by the partial lack of statistically differences of the methods in comparison to the expert pathologists, as shown in Fig 7. This is specially true for the scenario of pathology detection, where the participants even outperform one of the experts. However, this difference tends to reduce when considering abnormality detection. In fact, similarly to the automatic methods (recall Fig. 6b), the human observers, as shown in Fig. 8, have a high agreement level for Invasive cases but tend to disagree on the other classes. Figs. 7 and 8, together with Tables 4 and 5, suggest that specialists rely not only on objective markers, but also on their experience, intuition and personal assessment of cost of failure to perform the diagnosis, as assessable on the kappa score of the Pathological and Cancer-wise classifications. Also, similarly to the deep learning models, the specialists had more difficulty in distinguishing between the Normal and Benign classes in comparison with the cancerous classes. This further corroborates the hypothesis that the participants tended to fail Benign images due to the previously discussed complexity of this class.

Overall, the results in Tables 4 and 5, as well as the comparison of the quadratic Cohen's kappa score (Fig 7) between the different pathologists vs ground-truth and the automatic methods vs the ground-truth, show that deep learning models trained on properly annotated data can achieve human performance in complex medical tasks and may in the near future play an important role as second-opinion systems.

5.2. Performance in Part B

In general, Part B is much more challenging than Part A due to the large amount of information to process and need to integrate a wide range of scales. The pixel-wise sensitivity and specificity of the methods from Part B detailed in Tables 6 and 7, shows that the Invasive class tends to be the easiest to detect as the methods achieved an average sensitivity of 0.4. This is to be expected, since Invasive carcinoma is characterizable by an abnormal and non-confined nuclei density and thus methods tend to require less contextual information for the prediction. In fact, this is corroborated by the results in Part A that indicate that Invasive is the easiest of the pathological classes (see discussion of Fig. 6b). On the other hand, *In situ* has the lowest detection sensitivity (average of 0.06 and maximum value of 0.18) of the pathological classes. Unlike the Invasive carcinoma, classification as *in situ* is reliant on the location of the pathological cells – which means that without proper global and local context, which is complex to achieve due to the large size of the images, this class becomes non-trivial to classify. On the other hand, the methods of Part A did not tend to fail on images from the *in situ* class. This indicates that the microscopy images provide enough local and global context to perform the labeling and thus that human experience had an essential role during the acquisition and annotation of these images.

Fig. 9 shows examples of predictions from the top-3 performing participants. In general, one can observe the overestimation of invasive (blue) regions, and more difficulty in predicting the *in situ* (green) and benign (red) ones, which can also be seen in Table 6, where the sensitivity of the solutions for the invasive class is clearly superior to the others. Fig. 9 shows this tendency, where

two of the top performing teams fail to predict the *in situ* and benign regions and tend estimate them as invasive or as background.

5.3. Diversity in the solutions of the BACH challenge

Challenge designs should also promote a higher diversity of methodologies. However, BACH submissions followed the recent computer vision trend with deep learning vastly being the preferred approach. Specifically, pre-trained deep networks on natural images are relatively easy to set up and allow to achieve high performance while significantly reducing the need for field-knowledge, easing the participation on this and other challenges. Although the raw high performance of these methods is of interest, the scientific novelty of the approaches is reduced and usually limited to hyperparameter setup or network ensemble. Also, the black-box behavior of deep learning approaches hinders their application on the medical field, where specialists need to understand the reasoning behind the system's decision. It is the authors' belief that medical imaging challenges should further promote advances on the field by incentivating participants to propose significantly novel solutions that move from *what?* to *why?*. For instance, it would be of interest on future editions to ask participants to produce an automatic explanation of the method's decision. This will require the planning of new ground-truths and metrics that benefit systems that, by providing proper decision reasoning, are more capable of being used in the clinical practice.

5.3.1. Limitations of the BACH challenge

While an effort has been made in creating a relevant, stimulating and fair challenge, capable of advancing the state-of-the-art, the authors are aware of some limitations, namely: 1) The regional origin and relatively small size of the provided training dataset (specially for deep learning standards) may have limited the generalization capability of the solutions. Likewise, the relatively small test set does not allow to extensively evaluate the behaviour of the algorithms to different tissue structures and staining deviations. Increasing the diversity of the dataset would allow to draw even more relevant conclusions regarding the performance of image analysis systems for breast cancer diagnosis. 2) The reference labels for both Part **A** and Part **B** were obtained via manual annotation of two medical experts. Even though images where the observers disagreed were discarded, the labeling process is still reliant on the subjectivity and experience of the annotators (specially on Normal vs Benign labeling since no immunohistochemistry analysis is useful), limiting the performance of the submitted methods to that of the human expert. Increasing the number of annotators would allow to further increase the reliability of the dataset. 3) Patient-wise labels were only partially available for the training set. Participants could have used data with known patients for training and the remaining for method validation or use alternative approaches (such as clustering) to estimate the origin of the images. Despite this, the availability of this information for all images would allow a more fair patient-wise split for training and evaluating the algorithms, and could eventually lead to a smaller discrepancy between the training set cross-validated and the test set results. 4) The pixel-wise annotations of the WSI are not highly detailed and thus the delineated regions may include normal tissue regions besides the class assigned to that region. 5) Automatic evaluation of the participants' algorithms would have ease the submission procedure, allowing to an almost real-time feedback of the teams' performance. In this scenario, a scheme of multiple submissions could have been implemented, in which teams would be allowed to submit results on the website during the challenge running period, up to a maximum number of submissions. This would probably also boost the number of final submissions out of the challenge registrations.

6. Conclusions

BACH was organized to promote research on CAD systems for automatic breast cancer histology image analysis. Despite the complexity of the task, the challenge received a large number of high quality solutions that achieve similar performance to human experts. Namely, the best performing methods achieved 0.87 accuracy in classifying high resolution microscopy images in Normal, Benign, *In situ* carcinoma and Invasive carcinoma classes and a 0.69 score in labeling entire WSI.

Proper experiment design seems to be essential to achieve high performance in the breast cancer histology image analysis. Specifically, the conducted study allows to infer that 1) generically, using the latest CNN designs allows to positively impact the system's performance given that fine-tuning is properly performed; 2) CNNs seem to be robust to small color variations of H&E images and thus color normalization was not essential to attain high accuracies; 3) proper training splitting is essential to infer the generalization capability of the model, since CNNs may overfit to patient/acquisition details, and 4) using large context images as the network input allows for overall high performance even if the input image (and thus the overall quality of the information) has to be downsampled. On the other hand, current deep learning solutions still have issues dealing with large, high resolution images and further investment on development of methods for WSI analysis should be done.

It is the organizers' hope that the comprehensive analysis herein presented will motivate more challenges on medical imaging and specially pave the way for the development of new breast cancer CAD methods that contribute to the early detection of this pathology, with clear benefits for our societies.

Declaration of Competing Interest

The authors whose names are listed immediately below certify that they have NO affiliations with or involvement in any organization or entity with any financial interest (such as honoraria; educational grants; participation in speakers' bureaus; membership, employment, consultancies, stock ownership, or other equity interest; and expert testimony or patent-licensing arrangements), or non-financial interest (such as personal or professional relationships, affiliations, knowledge or beliefs) in the subject matter or materials discussed in this manuscript.

Author names:

Guilherme Aresta, Teresa Araújo, Scotty Kwok, Sai Saketh Chenamsetty, Mohammed Safwan, Varghese Alex, Bahram Marami, Marcel Prastawa, Monica Chan, Michael Donovan, Gerardo Fernandez, Jack Zeineh, Matthias Kohl, Christoph Walz, Florian Ludwig, Stefan Braunewell, Maximilian Baust, Quoc Dang Vu, Minh Nguyen Nhat To, Eal Kim, Jin Tae Kwak, Sameh Galal, Veronica Sanchez-Freire, Nadia Brancati, Maria Frucci, Daniel Riccio, Yaqi Wang, Lingling Sun, Kaiqiang Ma, Jiannan Fang, Ismael Kone, Lahsen Boulmane, Aurélio Campilho, Catarina Eloy, António Polónia, Paulo Aguiar

The authors whose names are listed immediately below report the following details of affiliation or involvement in an organization or entity with a financial or non-financial interest in the subject matter or materials discussed in this manuscript. Please specify the nature of the conflict on a separate sheet of paper if the space below is inadequate.

Acknowledgments

Guilherme Aresta is funded by the FCT grant contract [SFRH/BD/120435/2016](#). Teresa Araújo is funded by the FCT grant contract [SFRH/BD/122365/2016](#). Aurélio Campilho is with the

project "NanoSTIMA: Macro-to-Nano Human Sensing: Towards Integrated Multimodal Health Monitoring and Analytics/NORTE-0145-FEDER-000016", financed by the North Portugal Regional Operational Programme (NORTE 2020), under the PORTUGAL 2020 Partnership Agreement, and through the European Regional Development Fund (ERDF). Quoc Dang Vu, Minh Nguyen Nhat To, Eal Kim and Jin Tae Kwak are supported by the [National Research Foundation of Korea \(NRF\)](#) grant funded by the Korea government (MSIP) (No. 2016R1C1B2012433).

The authors would like to thank the pathologists Dr^a Ana Ribeiro, Dr^a Rita Canas Marques e Dr^a Ierecê Aymoré for their help in labeling the microscopy images.

The authors would also like to thank all the other BACH Challenge participants that registered, submitted their method and were accepted at the 15th International Conference on Image Analysis and Recognition (ICIAR 2018): Kamyar Nazeri, Azad Aminpour, and Mehran Ebrahimi; Nick Weiss, Henning Kost, and André Homeyer; Alexander Rakhlin, Alexey Shvets, Vladimir Iglovikov, and Alexandr A. Kalinin; Zeya Wang, Nanqing Dong, Wei Dai, Sean D. Rosario, and Eric P. Xing; Carlos A. Ferreira, Tânia Melo, Patrick Sousa, Maria Inês Meyer, Elham Shakibapour and Pedro Costa; Hongliu Cao, Simon Bernard, Laurent Heutte, and Robert Sabourin; Ruqayya Awan, Navid Alemi Koohbanani, Muhammad Shaban, Anna Lisowska, and Nasir Rajpoot; Sulaiman Vesal, Nishant Ravikumar, Amir Abbas Davari, Stephan Ellmann, and Andreas Maier; Yao Guo, Huihui Dong, Fangzhou Song, Chuang Zhu, and Jun Liu; Aditya Golatkar, Deepak Anand, and Amit Sethi; Tomas Iesmantas and Robertas Alzbutas; Wajahat Nawaz, Sagheer Ahmed, Ali Tahir, and Hassan Aqeel Khan; Artem Pimkin, Gleb Makarchuk, Vladimir Kondratenko, Maxim Pisov, Egor Krivov, and Mikhail Belyaev; Auxiliadora Sarmiento, and Irene Fondón; Quoc Dang Vu, Minh Nguyen Nhat To, Eal Kim, and Jin Tae Kwak; Yee-leng S. Vang, Zhen Chen, and Xiaohui Xie; Chao-Hui Huang, Jens Brodbeck, Nena M. Dimaano, John Kang, Belma Dogdas, Douglas Rollins, and Eric M. Gifford (authors are ordered according to conference proceedings).

References

- Abràmoff, M.D., Lou, Y., Erginay, A., Clarida, W., Amelon, R., Folk, J.C., Niemeijer, M., 2016. Improved automated detection of diabetic retinopathy on a publicly available dataset through integration of deep learning. In: *Investigative Ophthalmology & Visual Science*, 57, p. 5200. doi:10.1167/jovs.16-19964.
- [American Cancer Society](#), 2015. *Breast Cancer Facts & Figures 2015-2016*. American Cancer Society, Inc. Atlanta.
- Araújo, T., Aresta, G., Castro, E., Rouco, J., Aguiar, P., Eloy, C., Polónia, A., Campilho, A., 2017. Classification of breast cancer histology images using convolutional neural networks. *PLoS One* 12 (6), e0177544. doi:10.1371/journal.pone.0177544.
- Awan, R., Koohbanani, N.A., Shaban, M., Lisowska, A., Rajpoot, N., 2018. Context-aware learning using transferable features for classification of breast cancer histology images. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 788–795. doi:10.1007/978-3-319-93000-8_89.
- Bejnordi, B.E., Litjens, G., Timofeeva, N., Otte-Höller, I., Homeyer, A., Karssemeijer, N., Van Der Laak, J.A., 2016. Stain specific standardization of whole-slide histopathological images. *IEEE Trans. Med. Imaging* 35 (2), 404–415. doi:10.1109/TMI.2015.2476509.
- Bejnordi, B. E., Zuidhof, G., Balkenhol, M., Hermsen, M., Bult, P., van Ginneken, B., Karssemeijer, N., Litjens, G., van der Laak, J., 2017. Context-aware stacked convolutional neural networks for classification of breast carcinomas in whole-slide histopathology images. 1–13 arXiv:1705.03678. 10.1117/1.JMI.4.4.044504
- Belsare, A.D., Mushrif, M.M., Pangarkar, M.A., Meshram, N., 2015. Classification of breast cancer histopathology images using texture feature analysis. In: *TENCON 2015 - 2015 IEEE Region 10 Conference*, pp. 1–5. doi:10.1109/TENCON.2015.7372809.
- Brancati, N., Frucci, M., Riccio, D., 2018. Multi-classification of Breast cancer histology images by using a fine-tuning strategy. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 771–778. doi:10.1007/978-3-319-93000-8_87.
- Cao, H., Bernard, S., Heutte, L., Sabourin, R., 2018. Improve the performance of transfer learning without fine-tuning using dissimilarity-based multi-view learning for breast cancer histology images. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 779–787. doi:10.1007/978-3-319-93000-8_88.
- Chennamsetty, S.S., Safwan, M., Alex, V., 2018. Classification of breast cancer histology image using ensemble of pre-trained neural networks. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 804–811. doi:10.1007/978-3-319-93000-8_91.
- Cohen, J., 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* 20 (1), 37–46. doi:10.1177/001316446002000104.
- Cruz-Roa, A., Basavanahally, A., González, F., Gilmore, H., Feldman, M., Ganesan, S., Shih, N., Tomaszewski, J., Madabhushi, A., 2014. Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks. In: Gurcan, M.N., Madabhushi, A. (Eds.), *Medical Imaging 2014: Digital Pathology*, p. 904103. doi:10.1117/12.2043872. San Diego
- Cruz-Roa, A., Gilmore, H., Basavanahally, A., Feldman, M., Ganesan, S., Shih, N., Tomaszewski, J., Madabhushi, A., 2018. High-throughput adaptive sampling for whole-slide histopathology image analysis (HASHI) via convolutional neural networks: Application to invasive breast cancer detection(1), 1–23. 10.1371/journal.pone.0196828
- Cruz-Roa, A., Gilmore, H., Basavanahally, A., Feldman, M., Ganesan, S., Shih, N.N.C., Tomaszewski, J., González, F.A., 2017. Accurate and Reproducible Invasive Breast Cancer Detection in Whole- Slide Images: A Deep Learning Approach for Quantifying Tumor Extent. *Nature Publishing Group*, pp. 1–14. doi:10.1038/srep46450.
- Deng, J., Dong, W., Socher, R., Li-Jia Li, Kai Li, Fei-Fei, L., 2009. ImageNet: a large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. doi:10.1109/CVPRW.2009.5206848.
- Dieterich, T.G., 1993. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Comput.* 10 (7), 1895–1923.
- Edwards, A.L., 1948. Note on the "correction for continuity" In test—ing the significance of the difference between correlated proportions. *Psychometrik* 13 (3), 185–187.
- Elmore, J.G., Longton, G.M., Carney, P.A., Geller, B.M., Onega, T., Tosteson, A.N.A., Nelson, H.D., Pepe, M.S., Allison, K.H., Schnitt, S.J., O'Malley, F.P., Weaver, D.L., 2015. Diagnostic concordance among pathologists interpreting breast biopsy specimens. *JAMA* 313 (11), 1122. doi:10.1001/jama.2015.1405.
- Ferreira, C.A., Melo, T., Sousa, P., Meyer, M.I., Shakibapour, E., Costa, P., Campilho, A., 2018. Classification of breast cancer histology images through transfer learning using a pre-trained inception resnet V2. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 763–770. doi:10.1007/978-3-319-93000-8_86.
- Filipcuk, P., Fevens, T., Krzyzak, A., Monczak, R., 2013. Computer-aided breast cancer diagnosis based on the analysis of cytological images of fine needle biopsies. *IEEE Trans. Med. Imaging* 32 (12), 2169–2178. doi:10.1109/TMI.2013.2275151.
- Fondón, I., Sarmiento, A., García, A.I., Silvestre, M., Eloy, C., Polónia, A., Aguiar, P., 2018. Automatic classification of tissue malignancy for breast carcinoma diagnosis. *Comput. Biol. Med.* 96 (December 2017), 41–51. doi:10.1016/j.combiomed.2018.03.003.
- Galal, S., Sanchez-Freire, V., 2018. Candycane: breast cancer pixel-wise labeling with fully convolutional densenets. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 820–826. doi:10.1007/978-3-319-93000-8_93.
- Gecer, B., Aksoy, S., Mercan, E., Shapiro, L.G., Weaver, D.L., Elmore, J.G., 2018. Detection and classification of cancer in whole slide breast histopathology images using deep convolutional networks. *Pattern Recognit.* 84, 345–356. doi:10.1016/j.patcog.2018.07.022.
- George, Y.M., Zayed, H.H., Roushdy, M.I., Elbagoury, B.M., 2014. Remote computer-aided breast cancer detection and diagnosis system based on cytological images. *IEEE Syst. J.* 8 (3), 949–964. doi:10.1109/JYST.2013.2279415.
- Guo, Y., Dong, H., Song, F., Zhu, C., Liu, J., 2018. Breast cancer histology image classification based on deep neural networks. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 827–836. doi:10.1007/978-3-319-93000-8_94.
- Han, Z., Wei, B., Zheng, Y., Yin, Y., Li, K., Li, S., 2017. Breast cancer multi-classification from histopathological images with structured deep learning model. *Sci. Rep.* 7 (1), 1–10. doi:10.1038/s41598-017-04075-z.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778. doi:10.1109/CVPR.2016.90. arXiv:1512.03385.
- Hu, J., Shen, L., Sun, G., 2017. Squeeze-and-excitation networks. arXiv:1709.01507v2.
- Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, pp. 2261–2269. doi:10.1109/CVPR.2017.243.
- Iesmantas, T., Alzbutas, R., 2018. Convolutional capsule network for classification of breast cancer histology images. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), *Image Analysis and Recognition*. Springer International Publishing, Póvoa de Varzim, pp. 853–860. doi:10.1007/978-3-319-93000-8_97.
- Inoue, H., 2018. Data augmentation by pairing samples for images classification. arXiv:1801.02929.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: accelerating deep network training by reducing internal covariate shift. *Arxiv* 1–11. arXiv:1502.03167.
- Jégou, S., Drozdal, M., Vazquez, D., Romero, A., Bengio, Y., 2016. The one hundred layers tiramisú : fully convolutional densenets for semantic segmentation. arXiv:arXiv:1611.09326v3.
- Kingma, D.P., Ba, J., 2015. Adam: a method for stochastic optimization. In: *International Conference on Learning Representations (ICLR)*. San Diego
- Kohl, M., Walz, C., Ludwig, F., Brauneis, S., Baust, M., 2018. Assessment of breast cancer histology using densely connected convolutional networks. In:

- Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 903–913. doi:[10.1007/978-3-319-93000-8_103](https://doi.org/10.1007/978-3-319-93000-8_103).
- Koné, I., Boulmane, L., 2018. Hierarchical ResNeXt models for breast cancer histology image classification. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 796–803. doi:[10.1007/978-3-319-93000-8_90](https://doi.org/10.1007/978-3-319-93000-8_90).
- Kowal, M., Filipczuk, P., Obuchowicz, A., Korbicz, J., Monczak, R., 2013. Computer-aided diagnosis of breast cancer based on fine needle biopsy microscopic images. *Comput. Biol. Med.* 43 (10), 1563–1572. doi:[10.1016/j.combiomed.2013.08.003](https://doi.org/10.1016/j.combiomed.2013.08.003).
- Krishnan, M.M.R., Shah, P., 2012. Statistical analysis of textural features for improved classification of oral histopathological images. *J. Med. Syst.* 36 (2), 865–881. doi:[10.1007/s10916-010-9550-8](https://doi.org/10.1007/s10916-010-9550-8).
- Krizhevsky, A., Sutskever, I., Hinton, G., 2012. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* 25, 1106–1114.
- Kwok, S., 2018. Multiclass classification of breast cancer in whole-slide images. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 931–940. doi:[10.1007/978-3-319-93000-8_106](https://doi.org/10.1007/978-3-319-93000-8_106).
- Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. *Med. Image Anal.* 42, 60–88. doi:[10.1016/j.media.2017.07.005](https://doi.org/10.1016/j.media.2017.07.005).
- Litjens, G., Sánchez, C.I., Timofeeva, N., Hermesen, M., Nagtegaal, I., Kovacs, I., Hulsbergen-van de Kaa, C., Bult, P., van Ginneken, B., van der Laak, J., 2016. Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis. *Sci. Rep.* 6 (January), 26286. doi:[10.1038/srep26286](https://doi.org/10.1038/srep26286).
- Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Xiaojun Guan, Schmitt, C., Thomas, N.E., 2009. A method for normalizing histology slides for quantitative analysis. In: 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. IEEE, Boston, pp. 1107–1110. doi:[10.1109/ISBI.2009.5193250](https://doi.org/10.1109/ISBI.2009.5193250).
- Mahbod, A., Ellinger, I., Ecker, R., Smedby, Ö., Wang, C., 2018. Breast cancer histological image classification using fine-tuned deep network fusion. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 754–762. doi:[10.1007/978-3-319-93000-8_85](https://doi.org/10.1007/978-3-319-93000-8_85).
- Marami, B., Prastawa, M., Chan, M., Donovan, M., Fernandez, G., Zeineh, J., 2018. Ensemblenetwork for region identification in breast histopathology slides. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 861–868. doi:[10.1007/978-3-319-93000-8_98](https://doi.org/10.1007/978-3-319-93000-8_98).
- Pimkin, A., Makarchuk, G., Kondratenko, V., Pisov, M., Krivov, E., Belyaev, M., 2018. Ensembling Neural networks for digital pathology images classification and segmentation. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 877–886. doi:[10.1007/978-3-319-93000-8_100](https://doi.org/10.1007/978-3-319-93000-8_100).
- Rakhlina, A., Shvets, A., Igloukov, V., Kalinin, A.A., 2018. Deep convolutional neural networks for breast cancer histology image analysis. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 737–744. doi:[10.1007/978-3-319-93000-8_83](https://doi.org/10.1007/978-3-319-93000-8_83).
- Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P., 2001. Color transfer between images. *IEEE Comput. Graph. Appl.* 21 (4), 34–41. doi:[10.1109/38.946629](https://doi.org/10.1109/38.946629).
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI), 9351, pp. 234–241. doi:[10.1007/978-3-319-24574-4_28](https://doi.org/10.1007/978-3-319-24574-4_28).
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L., 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115 (3), 211–252. doi:[10.1007/s11263-015-0816-y](https://doi.org/10.1007/s11263-015-0816-y). arXiv:[1409.0575](https://arxiv.org/abs/1409.0575).
- Siegel, R.L., Miller, K.D., Jemal, A., 2017. Cancer statistics, 2017. *CA: Cancer J. Clin.* 67 (1), 7–30. doi:[10.3322/caac.21387](https://doi.org/10.3322/caac.21387).
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. 1–14 arXiv:[1409.1556](https://arxiv.org/abs/1409.1556).
- Smith, R.A., Cokkinides, V., Eyre, H.J., 2005. American cancer society guidelines for the early detection of cancer, 2006. *CA: Cancer J.Clin.* 55 (1), 31–44. doi:[10.3322/canjclin.53.1.27](https://doi.org/10.3322/canjclin.53.1.27).
- Spanhol, F.A., Oliveira, L.S., Petitjean, C., Heutte, L., 2016. A Dataset for Breast Cancer Histopathological Image Classification, 63, pp. 1455–1462.
- Spanhol, F.A., Oliveira, L.S., Petitjean, C., Heutte, L., 2016. Breast cancer histopathological image classification using Convolutional Neural Networks. In: 2016 International Joint Conference on Neural Networks (IJCNN). IEEE, Vancouver, pp. 2560–2567. doi:[10.1109/IJCNN.2016.7727519](https://doi.org/10.1109/IJCNN.2016.7727519).
- Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A., 2017. Inception-v4, inception-ResNet and the impact of residual connections on learning. In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17), pp. 4278–4284. arXiv:[1602.07261](https://arxiv.org/abs/1602.07261).
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 1–9. doi:[10.1109/CVPR.2015.7298594](https://doi.org/10.1109/CVPR.2015.7298594). arXiv:[1409.4842](https://arxiv.org/abs/1409.4842).
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2818–2826. doi:[10.1002/2014GB005021](https://doi.org/10.1002/2014GB005021). arXiv:[1512.00567](https://arxiv.org/abs/1512.00567).
- Tajbakhsh, N., Shin, J.Y., Gurudu, S.R., Hurst, R.T., Kendall, C.B., Gotway, M.B., Liang, J., 2016. Convolutional neural networks for medical image analysis: full training or fine tuning? *IEEE Trans. Med. Imaging* 35 (5), 1299–1312. doi:[10.1109/TMI.2016.2535302](https://doi.org/10.1109/TMI.2016.2535302).
- Vu, Q.D., To, M.N.N., Kim, E., Kwak, J.T., 2018. Micro and macro breast histology image analysis by partial network re-use. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 895–902. doi:[10.1007/978-3-319-93000-8_102](https://doi.org/10.1007/978-3-319-93000-8_102).
- Wang, Y., Sun, L., Ma, K., Fang, J., 2018. Breastcancer microscope image classification based on CNN with image deformation. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 845–852. doi:[10.1007/978-3-319-93000-8_96](https://doi.org/10.1007/978-3-319-93000-8_96).
- Wang, Z., Dong, N., Dai, W., Rosario, S.D., Xing, E.P., 2018. Classification of breast cancer histopathological images using convolutional neural networks with hierarchical loss and global pooling. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 745–753. doi:[10.1007/978-3-319-93000-8_84](https://doi.org/10.1007/978-3-319-93000-8_84).
- Weiss, N., Kost, H., Homeyer, A., 2018. Towards interactive breast tumor classification using transfer learning. In: Campilho, A., Karray, F., ter Haar Romeny, B. (Eds.), Image Analysis and Recognition. Springer International Publishing, Póvoa de Varzim, pp. 727–736. doi:[10.1007/978-3-319-93000-8_82](https://doi.org/10.1007/978-3-319-93000-8_82).
- Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K., 2017. Aggregated residual transformations for deep neural networks. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5987–5995. doi:[10.1109/CVPR.2017.634](https://doi.org/10.1109/CVPR.2017.634).
- Yu, F., Koltun, V., 2015. Multi-scale context aggregation by dilated convolutions. arXiv:[1511.07122v3](https://arxiv.org/abs/1511.07122v3).
- Zhang, B., 2011. Breast cancer diagnosis from biopsy images by serial fusion of random subspace ensembles. In: 2011 4th International Conference on Biomedical Engineering and Informatics (BMEI), 1, pp. 180–186. doi:[10.1109/BMEI.2011.6098229](https://doi.org/10.1109/BMEI.2011.6098229).