

Whole-Slide Mitosis Detection in H&E Breast Histology Using PHH3 as a Reference to Train Distilled Stain-Invariant Convolutional Networks

David Tellez¹, Maschenka Balkenhol, Irene Otte-Höller, Rob van de Loo, Rob Vogels, Peter Bult, Carla Wauters, Willem Vreuls, Suzanne Mol, Nico Karssemeijer, Geert Litjens, Jeroen van der Laak, and Francesco Ciampi

Abstract—Manual counting of mitotic tumor cells in tissue sections constitutes one of the strongest prognostic markers for breast cancer. This procedure, however, is time-consuming and error-prone. We developed a method to automatically detect mitotic figures in breast cancer tissue sections based on convolutional neural networks (CNNs). Application of CNNs to hematoxylin and eosin (H&E) stained histological tissue sections is hampered by noisy and expensive reference standards established by pathologists, lack of generalization due to staining variation across laboratories, and high computational requirements needed to process gigapixel whole-slide images (WSIs). In this paper, we present a method to train and evaluate CNNs to specifically solve these issues in the context of mitosis detection in breast cancer WSIs. First, by combining image analysis of mitotic activity in phosphohistone-H3 retained slides and registration, we built a reference standard for mitosis detection in entire H&E WSIs requiring minimal manual annotation effort. Second, we designed a data augmentation strategy that creates diverse and realistic H&E stain variations by modifying H&E color channels directly. Using it during training combined with network ensembling resulted in a stain invariant mitosis detector. Third, we applied knowledge distillation to reduce the computational requirements of the mitosis detection ensemble with a negligible loss of performance. The system was trained in a single-center cohort and evaluated in an independent multicenter cohort from the cancer genome atlas on the three tasks of the tumor proliferation assessment challenge. We obtained a performance within the top three best methods for most of the tasks of the challenge.

Index Terms—Breast cancer, mitosis detection, convolutional neural networks, phosphohistone-H3, data augmentation, knowledge distillation.

I. INTRODUCTION

HISTOPATHOLOGICAL tumor grade is a strong prognostic marker for the survival of breast cancer patients [1], [2]. It is assessed by examination of hematoxylin and eosin (H&E) stained tissue sections using bright-field microscopy [3]. Histopathological grading of breast cancer combines information from three morphological features: (1) nuclear pleomorphism, (2) tubule formation and (3) mitotic count, and can take a value within the 1-3 range, where 3 corresponds to the worst patient prognosis. In this study, we focus our attention on the mitosis count component, as it can be used as a reliable and independent prognostic marker [2], [4]. Mitosis is a crucial phase in the cell cycle where a replicated set of chromosomes is split into two individual cell nuclei. These chromosomes can be recognized in H&E stained sections as mitotic figures (see Fig. 1). For breast cancer grading, the counting of mitotic figures is performed by first identifying a region of 2 mm² with a high number of mitotic figures at low microscope magnification (hotspot) and subsequently counting all mitotic figures in this region at high magnification.

The recent introduction of whole-slide scanners in anatomic pathology enables pathologists to make their diagnoses on digitized slides [5], so-called whole-slide images (WSIs), and promotes the development of novel image analysis tools for automated and reproducible mitosis counting. Publicly available training datasets for mitosis detection [6]–[9] have important limitations in terms of (1) size, (2) tissue representativity, and (3) reference standard agreement. In these datasets, the total number of annotated mitotic figures is currently limited to 1500 objects, far away from standard datasets used to train modern computer vision systems [10], [11]. In addition, mitotic figures were annotated in certain manually selected tumor regions only, often excluding tissue areas with image artifacts (common in WSIs). Furthermore, exhaustive manual annotations are known to suffer from disagreement among observers and limited recall [12], [13]. We propose a method to improve the annotation process based on the automatic analysis of immunohistochemical stained slides.

Manuscript received February 1, 2018; revised March 15, 2018; accepted March 20, 2018. Date of publication March 28, 2018; date of current version August 30, 2018. This work was supported by the Radboud Institute of Health Sciences, Nijmegen, The Netherlands. (Corresponding author: David Tellez.)

D. Tellez, M. Balkenhol, I. Otte-Höller, R. van de Loo, G. Litjens, J. van der Laak, and F. Ciampi are with the Diagnostic Image Analysis Group, Radboud University Medical Center, 6500HB Nijmegen, The Netherlands, and also with the Department of Pathology, Radboud University Medical Center, 6500HB Nijmegen, The Netherlands (e-mail: david.tellezmartin@radboudumc.nl).

R. Vogels and P. Bult are with the Department of Pathology, Radboud University Medical Center, 6500HB Nijmegen, The Netherlands.

C. Wauters and W. Vreuls are with the Department of Pathology, Canisius-Wilhelmina Hospital, 6532SZ Nijmegen, The Netherlands.

S. Mol is with the Department of Pathology, Jeroen Bosch Hospital, 5223GZ 's-Hertogenbosch, The Netherlands.

N. Karssemeijer is with the Diagnostic Image Analysis Group, Radboud University Medical Center, 6500HB Nijmegen, The Netherlands.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMI.2018.2820199

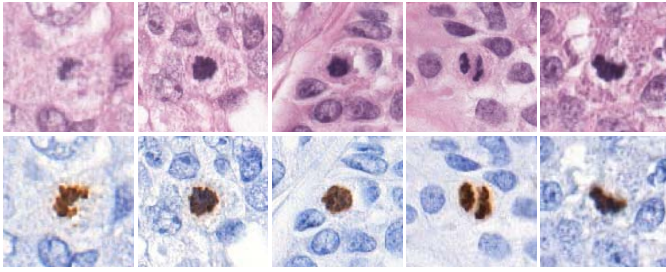


Fig. 1. Examples of image patches containing mitotic figures, shown at the center of each patch. In H&E (top), mitotic figures are visible as dark spots. In PHH3 (bottom), they are visible as brown spots. Mitotic figures in PHH3 stain are easier to identify than in H&E stain.

Phosphohistone-H3 (PHH3) is an antibody that identifies cells undergoing mitosis [14], [15]. Mitotic figures appear in PHH3 immunohistochemically stained slides (abbreviated as PHH3 stained slides) as high contrast objects that are easier to detect than in H&E [16]–[18], illustrated in Fig. 1. We propose to *destain* H&E slides and *restain* them with PHH3 to obtain both H&E and PHH3 WSIs from the exact same tissue section [19]. By automatically analyzing mitotic activity in PHH3 and registering it to H&E, we generated training data for mitosis detection in H&E WSIs in a scalable manner, i.e. independent from the manual annotation procedure.

Although the process of H&E tissue staining follows a standard protocol, the appearance of the stained slides is not identical among pathology laboratories and varies across time even within the same center (see Fig. 2). This variance typically causes mitosis detection algorithms to underperform on images originating from pathology laboratories different than the one that provided the training data [13]. Several solutions have been proposed to tackle this lack of generalization. First, building multicenter training datasets that contain sufficient stain variability. Following this approach, the Tumor Proliferation Assessment Challenge (TUPAC) resulted in numerous successful mitosis detection algorithms. Top-performing methods in the challenge are based on convolutional neural networks (CNNs) [12], [13], [20]–[22]. This is in line with the trend observed in recent years, which has seen CNNs as the top-performing approach in image analysis, both in computer vision and medical imaging [23], and corroborates the fact that CNNs have become the standard methodology for automatic mitosis detection. However, multicenter datasets cannot cover all the variability encountered in clinical practice, and are expensive to collect. Second, stain standardization techniques [24]–[26] have been widely used by many of these successful mitosis detection methods to reduce stain variability. However, they require preprocessing all training and testing WSIs and do not reduce the generalization error of trained models. Third, data augmentation strategies have been used to simulate stain variability during the model training. These techniques typically involve RGB transformations such as brightness and contrast enhancements, and color hue perturbations [13], [27]. We argue that designing specific data augmentation strategies for H&E stained tissue images is the most promising approach to reduce the generalization error of these networks, avoiding the elevated costs of assembling a

multicenter cohort, and effectively enforcing stain invariance into the trained models. We propose an augmentation strategy tailored to H&E WSIs that modifies the hematoxylin and eosin color channels directly, as opposed to RGB, and it is able to generate a broad range of realistic H&E stain variations from images originating in a single center. We call this technique *stain augmentation*.

Automatic mitosis detection algorithms rely on techniques such as the use of high capacity CNNs and multi-network ensembling to achieve state of the art performance [9], [12], [21]. These are simple yet effective mechanisms to improve performance, reduce generalization error and diminish the sensitivity of the model to the detection threshold. However, due to their computationally expensive nature, it is unfeasible to use them for dense prediction in gigapixel WSIs. We propose to exploit the technique of *knowledge distillation* [28] to reduce the size of the trained ensemble to that of a single network, maintaining similar levels of performance and increasing processing speed drastically.

Our contributions can be summarized as follows:

- We propose a scalable procedure to exhaustively annotate mitotic figures in H&E WSIs with minimal human labeling effort. We do so by automatically analyzing mitotic activity in PHH3 restained tissue sections and registering it to H&E.
- We propose a data augmentation technique that generates a broad range of realistic H&E stain variations by modifying the hematoxylin and eosin color channels directly. We demonstrate its ability to enforce stain invariance by transferring the performance obtained in a dataset from a single center to a multicenter publicly available cohort.
- We apply knowledge distillation to reduce the size of an ensemble of trained networks to that of a single network, in order to perform mitosis detection in gigapixel WSIs with similar performance and vastly increased processing speed.

The paper is organized as follows. Sec. II reports the datasets used to train and validate our method. Sec. III and Sec. IV describe the methodology in depth. All details regarding CNN architectures, training protocols and hyperparameter tuning are explained in Sec. V. Experimental results are listed in Sec. VI, followed by Sec. VII where the discussion and final conclusion are stated.

II. MATERIALS

In this study, we use two cohorts from different sources for (1) developing the main mitosis detection algorithm, and (2) performing an independent evaluation of the system performance. Details on the datasets are provided in Table I.

The first cohort consists of 18 triple negative breast cancer (TNBC) patients who underwent surgery in three different hospitals in the Netherlands: Jeroen Bosch Hospital, Radboud University Medical Centre (Radboudumc) and Canisius-Wilhelmina Hospital. However, all tissue sections were cut, H&E stained and scanned at the Radboudumc using a 3DHis-tech Panoramic 250 Flash II scanner at a spatial resolution of $0.25\mu\text{m}/\text{pixel}$, therefore, we consider this set of WSIs a single-center one. Subsequently, the slides were destained,

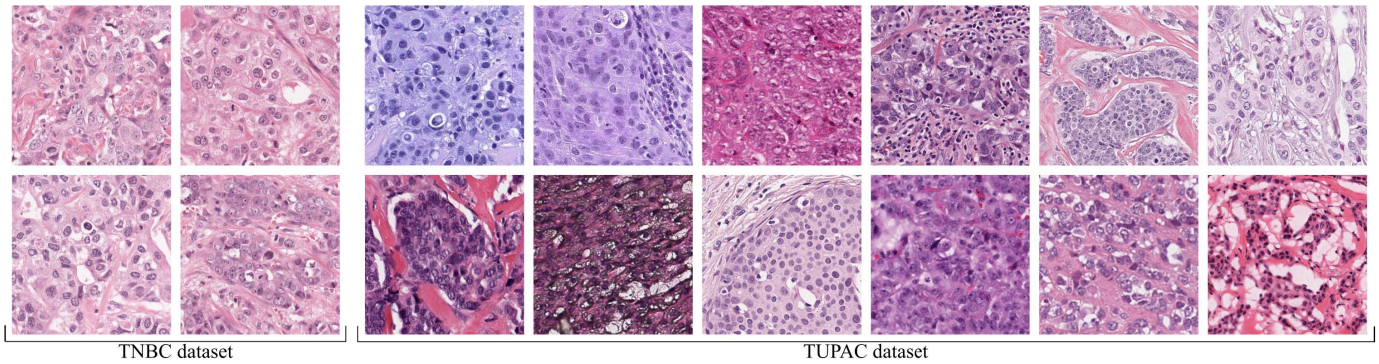


Fig. 2. Breast tissue samples stained with hematoxylin and eosin (H&E). Each tile comes from a different patient. The triple negative breast cancer (TNBC) cohort contains images from a single center, whereas the Tumor Proliferation Assessment Challenge (TUPAC) dataset contains images from multiple centers. Notice the homogeneous appearance of the *TNBC* dataset, used for training our mitosis detector, and the variable stain of the *TUPAC* dataset, used to validate our method.

TABLE I

OVERVIEW OF THE DATASETS USED IN THIS STUDY. THE PURPOSE OF *Training* INDICATES THAT THE DATASET WAS USED TO TRAIN A CNN FOR MITOSIS DETECTION, WHEREAS *Threshold Tuning* INDICATES THAT IT WAS SOLELY EMPLOYED TO OPTIMIZE CERTAIN HYPER-PARAMETERS, SUCH AS THE DETECTION THRESHOLD

Alias	Data	Cases	Multicenter	Reference standard	Purpose
TNBC-H&E	H&E whole-slide images	18	No	None	Training
TNBC-PHH3	PHH3 whole-slide images	18	No	Set of annotated patches	Training
TUPAC-train	H&E whole-slide images	493	Yes	Grading and proliferation score	Threshold tuning
TUPAC-aux-train	H&E selected regions	50	Yes	Exhaustive location of mitotic figures	Threshold tuning
TUPAC-test	H&E whole-slide images	321	Yes	Not publicly available	Evaluation
TUPAC-aux-test	H&E selected regions	34	Yes	Not publicly available	Evaluation

restained with PHH3 and re-scanned, resulting in 18 pairs of H&E and PHH3 WSIs representing the exact same tissue section per pair. We will refer to these images as the *TNBC-H&E* and *TNBC-PHH3* datasets through the rest of the paper.

The second cohort consists of the publicly available TUPAC dataset [9]. In particular, 814 H&E WSIs from invasive breast cancer patients from multiple centers included in The Cancer Genome Atlas [29] scanned at $0.25\mu\text{m}/\text{pixel}$ were annotated, providing two labels for each case. The first label is the histological grading of each tumor based on the mitotic count only. The second score is the outcome of a molecular test highly correlated with tumor proliferation [30]. Out of these 814 WSIs, 493 cases have a public reference standard, whereas the remaining 321 cases do not (only available to TUPAC organizers). We will refer to these sets of WSIs as the *TUPAC-train* and *TUPAC-test* datasets respectively.

Additionally, the organizers of TUPAC provided data for individual mitotic figure detection from two centers. They preselected 84 proliferative tumor regions from H&E WSIs of breast cancer cases, and two observers exhaustively annotated them. Coincident annotations were marked as the reference standard, and discording items were reviewed by a third pathologist. Out of these 84 regions, 50 cases have a public reference standard, whereas the remaining 34 cases do not (only available to TUPAC organizers). We will refer to these sets as the *TUPAC-aux-train* and *TUPAC-aux-test* datasets respectively.

All WSIs in this work were preprocessed with a tissue-background segmentation algorithm [31] in order to exclude areas not containing tissue from the analysis.

III. PHH3 STAIN: REFERENCE STANDARD FOR MITOTIC ACTIVITY

We trained a CNN to automatically identify mitotic figures throughout PHH3 WSIs and registered the detections to the H&E WSI pairs. The entire process is summarized in Fig. 3.

A. Mitosis Detection in PHH3 Whole-Slide Images

We used color deconvolution to disentangle the DAB and hematoxylin stains (brown and blue color channels respectively) and obtained the mitotic candidates by labeling connected components of all positive pixels belonging to the DAB stain. We observed that the PHH3 antibody was sensitive but not specific regarding mitosis activity: 75% of candidates were trivial artifacts. To avoid waste of manual annotation effort on these artifacts, we trained a CNN, named *CNN1*, to classify candidates among artifactual or non-artifactual objects. To train such a system, a student labeled 2000 randomly selected candidates. We classified all PHH3 candidates with *CNN1*, and randomly selected 2000 samples classified as non-artifactual. Four observers labeled these samples as either containing a mitotic figure or not. These four observers consisted of a pathology resident, a PhD student and two lab technicians, and they were provided with sufficient training, visual examples and hands-on practice on mitosis detection. Annotations were aggregated by performing majority voting, keeping only those samples where at least 3 observers agreed upon. This resulted in 778 and 1093 mitotic and non-mitotic annotations, respectively. Furthermore, the non-mitotic set was extended with 1500 artifactual samples used to train *CNN1*, resulting in 2593 non-mitotic samples.

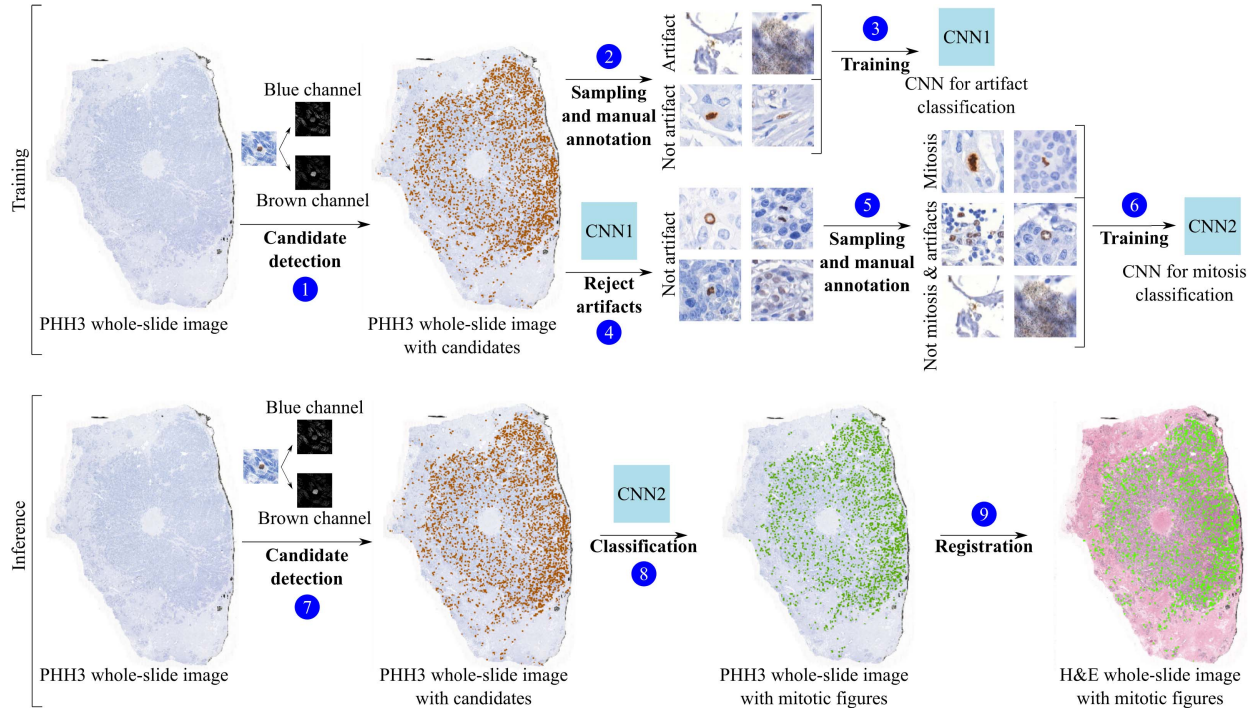


Fig. 3. Building reference standard for mitotic activity using PHH3 stained slides. Top: training stage, where mitotic candidates are extracted from the brown color channel (1), pruned from artifacts (2, 3, 4), a subset is manually annotated (5), and then used to train a CNN to distinguish between mitotic and non-mitotic patches (6), named *CNN2*. Bottom: inference stage, where candidates in a given PHH3 slide (7) are classified with *CNN2* as mitotic or non-mitotic (8), then registered to their respective H&E slide pairs (9).

We used this annotated set of samples to train *CNN2* to distinguish PHH3 candidates among mitotic and non-mitotic patches. During training, we randomly applied several techniques to augment the data and prevent overfitting, namely: rotations, vertical and horizontal mirroring, elastic deformation [32], Gaussian blurring, and translations. The resulting performance of *CNN2* was an F1-score of 0.9. Details on network architecture, training protocol and hyper-parameter selection are provided in Sec. V. We classified all candidates found in the PHH3 slides as mitotic or non-mitotic objects using *CNN2*, generating exhaustive reference standard data for mitosis activity at whole-slide level.

B. Registering Mitotic Figures From PHH3 to H&E Slides

The process of slide restaining guaranteed that the exact same tissue section was present in both the H&E and the PHH3 WSIs, requiring minimal registration to align mitotic objects. We designed a simple yet effective two-step routine to reduce vertical and horizontal shift at *global* and *local* scale. First, we performed a global and coarse registration that minimized the vertical and horizontal shift between image pairs. We did so by finding the alignment of randomly selected pairs of corresponding tiles as the shift vector that maximized the 2D cross-correlation. For improved accuracy, we repeated this procedure 10 times per WSI pair (at random locations throughout the WSI), averaging the cross-correlation heatmap across trials. Finally, all mitoses were adjusted with the resulting global shift vector. Second, we registered each mitotic figure individually, following a similar procedure as before. We extracted a single pair of high magnification tiles, centered in each candidate location, to account for individual local shifts.

IV. H&E STAIN: TRAINING A MITOSIS DETECTOR

We trained a CNN for the task of mitosis detection and used it to exhaustively locate mitotic figures throughout H&E WSIs. Only slides from the *TNBC-H&E* dataset were used in this step. The procedure is summarized in Fig. 4.

A. Assembling a Training Dataset for Mitosis Detection

Even though the *TNBC-H&E* dataset already possessed little stain variation, we standardized the stain of each WSI to reduce intra-laboratory stain variations [26], preventing the CNN from becoming stain invariant from the raw training data. This further strengthens the challenge of generalizing to unseen stain variations.

As a result of the large amount of available pixels in WSIs, the selection of negative samples was not trivial. We propose a *candidate detector* based on the assumption that mitotic figures are non-overlapping circular objects with dark inner cores. This detector found candidates by iterative detection and suppression of all circular areas of diameter d centered on local minima of the mean RGB intensity until all pixels above a threshold t were exhausted. We selected a sufficiently large t so that candidates were representative for all tissue types. A candidate was labeled as a positive sample if its Euclidean distance to any reference standard object was at most d pixels, and labeled as a negative sample otherwise.

Most of the negative samples were very easy to classify and their contribution to improve the decision boundaries of the CNN was marginal. We found it crucial to identify highly informative negative samples to train the CNN effectively. We proceeded similarly as stated in [12]. First, we built an *easy* training dataset by including all positive candidates and a

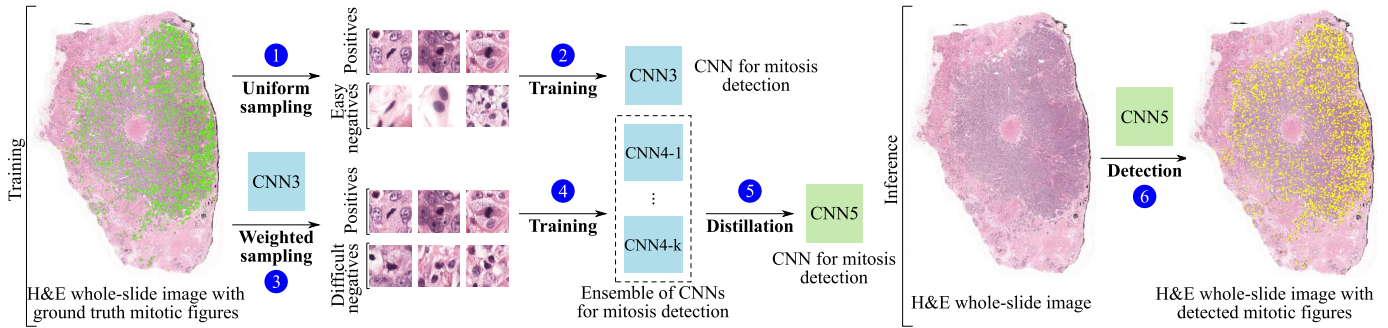


Fig. 4. Training a whole-slide H&E mitosis detector. Left (training stage): an auxiliary network *CNN3* is trained with uniformly sampled patches (1, 2). Then, *CNN3* is used to perform hard mining on the negative patches and create k distinct datasets sampled via bootstrapping (3). These datasets are used to train networks *CNN4-1* to *CNN4-k* and build an ensemble (4). Finally, the ensemble is reduced into *CNN5* via knowledge distillation (5). Right (inference stage): *CNN5* is applied in a fully convolutional manner throughout an entire WSI to detect mitotic figures exhaustively (6).

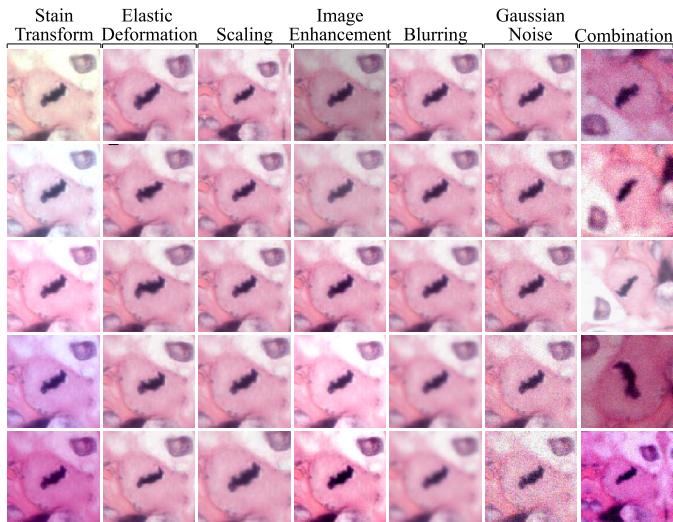


Fig. 5. Multiple augmented versions of the same mitotic patch. Each column shows samples of a single augmentation function except for the last one, which combines all the techniques together with rotation and mirroring.

number of uniformly sampled negative candidates, and trained a network with it, labeled as *CNN3* for future reference. Second, we evaluated all candidate patches with this network, obtaining a prediction probability for each of them. Finally, we built a *difficult* training dataset by selecting all positive candidates, and a number of negative candidates sampled proportionally to their probability of being mitosis, so that harder samples were chosen more often.

B. H&E Stain Augmentation

We trained a CNN on the *difficult* dataset to effectively distinguish between mitotic and non-mitotic H&E patches, named as *CNN4*. During training, we applied several techniques to augment the dataset on-the-fly, preventing overfitting and improving generalization. We implemented several routines for H&E histopathology imaging in the context of mitosis detection, illustrated in Figure 5 with a sample patch.

1) *Morphology Invariance*: We exploited the fact that mitotic figures can have variable shapes and sizes by augmenting the training patches with rotation, vertical and horizontal mirroring (R), scaling (S), elastic deformation (E) [32], and translation around the central pixel.

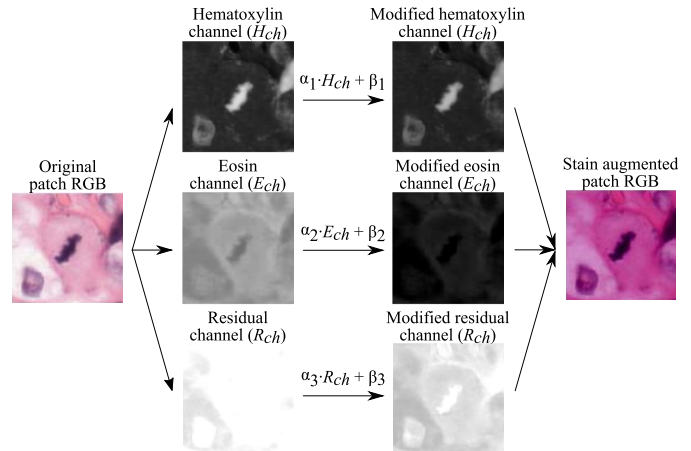


Fig. 6. H&E stain augmentation. From left to right: first, an RGB patch is decomposed into hematoxylin (H_{ch}), eosin (E_{ch}) and residual (R_{ch}) color channels. Then, each channel is individually modified with a random factor and bias. Finally, resulting channels are transformed back to RGB color space.

2) *Stain Invariance*: We used a novel approach to simulate a broad range of realistic H&E stained images by retrieving and modifying the intensity of the hematoxylin and eosin color channels directly (C), as illustrated in Figure 6. First, we transformed each patch sample from RGB to H&E color space using a method based on color deconvolution [33], see the Appendix for more methodological details. Second, we modified each channel i individually, i.e. hematoxylin (H_{ch}), eosin (E_{ch}) and residual (R_{ch}), with random factors α_i and biases β_i taken from two uniform distributions. Finally, we combined and projected the resulting color channels back to RGB. Additionally, we simulated further alternative stains by modifying image brightness, eosness, contrast and color intensity (H).

3) *Artifact Invariance*: We mimicked the out of focus effect of whole-slide scanners with a Gaussian filter (B), and added Gaussian noise to decrease the signal-to-noise ratio of the images (G), simulating image compression artifacts.

C. Ensemble & Network Distillation

The use of an ensemble of networks is a key factor to achieve state of the art performance in multiple classification tasks [11], particularly in mitosis detection [12], [21]. An ensemble of networks performs significantly better than

TABLE II

ARCHITECTURE OF *CNN3*, *CNN4* AND *CNN5*. γ CONTROLS THE NUMBER OF FILTERS PER CONVOLUTIONAL LAYER

Function	Filters	Size	Stride	Activation
conv	32γ	3x3	1	Leaky-ReLU
conv	32γ	3x3	2	Leaky-ReLU
conv	64γ	3x3	1	Leaky-ReLU
conv	64γ	3x3	2	Leaky-ReLU
conv	128γ	3x3	1	Leaky-ReLU
conv	128γ	3x3	1	Leaky-ReLU
conv	256γ	3x3	1	Leaky-ReLU
conv	256γ	3x3	1	Leaky-ReLU
conv	512γ	14x14	1	Leaky-ReLU
dropout	-	-	-	-
conv	2	1	1	Softmax

any of its members if they make independent (uncorrelated) errors. Building different training datasets with replacement (bagging) has been shown to increase model independence in an ensemble [20]. Therefore, we trained k different CNNs on k different training datasets, each one obtained by sampling negative candidates with replacement, and made an ensemble with the networks, averaging their predicted probabilities across models.

The computational requirements of this ensemble grow proportionally to k . To reduce this burden, we applied the idea of knowledge distillation, a technique designed to transfer the performance of a network (or ensemble of networks) to a lighter target neural network [28]. To achieve the highest performance, we distilled the ensemble of k networks to a single smaller CNN, named *CNN5*. We did so by training *CNN5* directly on the continuous averaged output probabilities of the ensemble, instead of the dataset labels, as indicated in [28]. We defined γ as a parameter to control the amount of trainable parameters used by *CNN5*, taking values in the $[0, 1]$ range. In particular, the number of filters per convolutional layer was proportional to this parameter. It defaults to $\gamma = 1$, unless stated otherwise.

D. Outcome at Whole-Slide Level

We detected mitotic figures at whole-slide level by sliding *CNN5* over tissue regions at $0.25\mu\text{m}/\text{pixel}$ resolution, producing a mitosis probability map image for each slide. Simple post-processing allowed to detect individual mitotic objects: (1) thresholding and binarizing the probability map (pixel probability of at least 0.8 to reduce the computational burden), (2) labeling connected components in the resulting image, and (3) suppressing double detections closer than d pixels. A detection probability threshold δ must be provided to discern between false and true positives.

In order to provide the number of mitotic figures in the most active tumor area per slide, we slid a virtual hotspot consisting of a 2 mm^2 circle throughout each entire WSI, counting the number of mitoses at each unique spatial position. To identify the hotspot with the largest mitotic activity ignoring large outliers, we considered the 95th percentile of the series of mitotic countings for each slide, excluding empty areas.

To estimate the tumor grading of the patient, we followed the guidelines used to compute the Mitotic Activity Index (MAI) [3]. We defined two thresholds, θ_1 and θ_2 , and used

them to categorize the number of mitotic figures in the hotspot into three possible outcomes. In particular, we predicted grade 1, 2 or 3 depending on whether the mitotic counting was below or equal to θ_1 , between thresholds, or above θ_2 , respectively. To estimate a continuous tumor proliferation score, we simply provided the number of mitotic figures in the hotspot.

V. CNN ARCHITECTURE, TRAINING PROTOCOL AND OTHER HYPER-PARAMETERS

In order to train the CNN models, we used RGB patches of 128×128 pixel size, taken at $0.25\mu\text{m}/\text{pixel}$ resolution and whose central pixel was centered in the coordinates of the annotated object. Patches were cropped as part of the data augmentation strategy, resulting in 100×100 pixel images fed into the CNNs.

Convolutional neural networks were trained to minimize the cross-entropy loss of the network outputs with respect to the patch labels, using stochastic gradient descent with Adam optimization and balanced mini-batch of 64 samples. To prevent overfitting, an additional L_2 term was added to the network loss, with a factor of $1e-5$. Furthermore, the learning rate was exponentially decreased from $1e-3$ to $3e-5$ through 20 epochs. At the end of training, network parameters corresponding to the checkpoint with the highest validation F1-score were selected for inference.

A. Mitosis Detection in PHH3

CNN1 was trained with 1500 artifactual and 500 non-artifactual samples, and *CNN2* was trained with 778 mitotic and 2593 non-mitotic patch samples. In both cases, the sets of samples were randomly divided into training and validation subsets at case level from *TNBC-PHH3*, with 10 and 8 slides for training and validation, respectively. The architecture of *CNN1* and *CNN2* consisted of five pairs of convolutional and max-pooling layers with 64, 128, 256, 512 and 1024 3×3 filters per layer, followed by two densely connected layers of 2048 and 2 units respectively. A dropout layer was placed between the last two layers, with 0.5 coefficient. All convolutional and dense layers were followed by ReLU functions, except for the last layer that ended with a softmax function.

B. Mitosis Detection in H&E

The H&E slides provided in the *TNBC-H&E* dataset were randomly divided into training, validation and test subsets, with 11, 3 and 4 slides each. For the candidate detector, we selected $d = 100$ as the diameter of an average tumor cell at $0.25\mu\text{m}/\text{pixel}$ resolution; and $t = 0.6$ recalling 99% of the mitotic figures in the reference standard of the validation set, with a rate of 1 positive to every 1000 negative samples. On average, each slide had 1 million candidates. The architecture of *CNN3*, *CNN4* and *CNN5* is summarized in Table II. We found that substituting max-pooling layers for strided convolutions slightly improved convergence, resulting in an *all convolutional* architecture [34]. To train *CNN3*, we built a training set with all positive samples, 22,000 mitotic figures, and 100,000 uniformly sampled negative candidates. For validation purposes, we also built a validation set consisting

TABLE III

ANALYSIS OF THE IMPACT IN PERFORMANCE OF USING DATA AUGMENTATION, ENSEMBLE AND KNOWLEDGE DISTILLATION. EACH ROW REPRESENTS A CNN (OR SET OF CNNs FOR THE ENSEMBLE CASE) THAT WAS TRAINED USING *TNBC-H&E* DATA AS EXPLAINED IN SEC. IV. EACH TRAINED NETWORK WAS EVALUATED IN THE INDEPENDENT *TUPAC-aux-train* DATASET WITH MULTIPLE DETECTION THRESHOLDS, REPORTING THE HIGHEST F1-SCORE OBTAINED AND THE NUMBER OF TRAINABLE PARAMETERS. EXPERIMENTS 1, 2 AND 3 COMPARED THE USE OF DIFFERENT DATA AUGMENTATION STRATEGIES (R: ROTATION, C: COLOR STAIN, S: SCALING, ETC., SEE SEC. IV.B FOR THE FULL LIST). EXPERIMENT 4 SHOWED THE PERFORMANCE OF AN ENSEMBLE OF $k = 10$ NETWORKS, TRAINED AS EXPLAINED IN SEC. IV.C. IN EXPERIMENTS 5, 6 AND 7, THE ENSEMBLE OF CNNs (EXPERIMENT 4) WAS DISTILLED INTO SINGLE SMALLER CNNs WITH VARYING CAPACITIES $\gamma = 1.0, 0.8, 0.6$ AS EXPLAINED IN SEC. IV.C. THE CNN TRAINED FOR EXPERIMENT 7 COINCIDES WITH *CNN5*.

Exp	Augment	Ensemble	Distilled	F1-score	Param
1	RSEB	No	No	0.018	26.9M
2	RCSEB	No	No	0.412	26.9M
3	RCSEHBG	No	No	0.613	26.9M
4	RCSEHBG	$k = 10$	No	0.660	269M
5	RCSEHBG	No	$\gamma = 1.0$	0.623	26.9M
6	RCSEHBG	No	$\gamma = 0.8$	0.628	17.1M
7	RCSEHBG	No	$\gamma = 0.6$	0.636	9.5M

of 10% of the total available samples in both classes, 200,000 negative and 500 positive candidates. To train *CNN4* and *CNN5*, we built a training set with all positive samples, 22,000 mitotic figures, and 400,000 negative candidates, sampling difficult patches more often with replacement. We used the same validation set as with *CNN3*. For the ensemble, we selected $k = 10$, the highest number of networks that we could manage in an ensemble with our computational resources. We distilled several versions of *CNN5* varying the value of γ , evaluated each network in the validation set of the *TNBC-H&E* dataset, and obtained an F1-score of 0.634, 0.646 and 0.629 for γ values of 1.0, 0.8 and 0.6, respectively. We selected $\gamma = 0.6$ for further experiments, resulting in a distilled network with 28X and 2.8X times less parameters than the ensemble and the single network ($\gamma = 1.0$), respectively, at a negligible cost of performance.

The color augmentation technique sampled α and β from two uniform distributions with ranges $[0.95, 1.05]$ and $[-0.05, 0.05]$, respectively. Patches were scaled with a zooming factor uniformly sampled from $[0.75, 1.25]$. The elastic deformation routine used $\alpha = 100$, and $\sigma = 10$. Color, contrast and brightness intensity was enhanced by factors uniformly sampled from $[0.75, 1.5]$, $[0.75, 1.5]$ and $[0.75, 1.25]$, respectively. The Gaussian filter used for blurring sampled σ uniformly from the $[0, 2]$ range. The additive Gaussian noise had zero mean and a standard deviation uniformly sampled from the $[0, 0.1]$ range. These parameters were selected empirically to simultaneously maximize visual variety and result in realistic samples.

VI. EXPERIMENTS AND RESULTS

A. Impact of Augmentation, Ensembling and Distillation

We performed a series of experiments to quantitatively assess the impact in performance of three ideas mentioned in

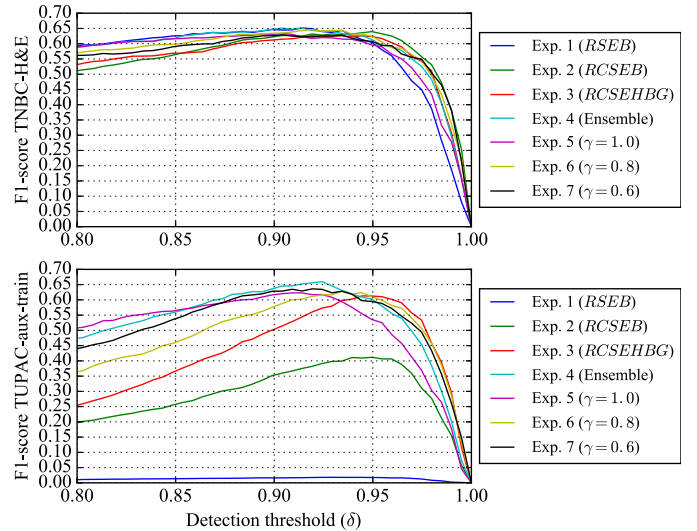


Fig. 7. Analysis of the impact in performance of using data augmentation, ensemble and knowledge distillation measured in terms of F1-score with respect to the detection threshold (top: *TNBC-H&E* dataset, bottom: *TUPAC-aux-train* dataset). Experiments 1, 2 and 3 compared the use of different data augmentation strategies (R: rotation, C: color stain, S: scaling, etc., see Sec. IV.B for the full list). Experiment 4 showed the performance of an ensemble of $k = 10$ networks, trained as explained in Sec. IV.C. In experiments 5, 6 and 7, the ensemble of CNNs (experiment 4) was distilled into single smaller CNNs with varying capacities $\gamma = 1.0, 0.8, 0.6$ as explained in Sec. IV.C. The CNN trained for experiment 7 coincides with *CNN5*.

this paper: (a) data augmentation, (b) ensemble and (c) knowledge distillation. In each experiment, we trained a CNN (or set of CNNs for the ensemble case) with the *TNBC-H&E* dataset as explained in Sec. IV and V. Each trained model was evaluated in the independent *TUPAC-aux-train* dataset with multiple detection thresholds, reporting the highest F1-score obtained. Table III summarizes the numerical results, and Figure 7 analyzes the sensitivity of each model with respect to the detection threshold.

1) *Data Augmentation*: The goal of experiments 1, 2 and 3 is to test whether the proposed data augmentation strategy can improve the performance of the CNN in an independent test set, in particular the novel color stain augmentation. In experiment 1, we trained a baseline system including only basic augmentation (*RSEB*) and obtained an F1-score of 0.018. In experiment 2, we repeated the training procedure including our color augmentation technique as well (*RCSEB*) and obtained an F1-score of 0.412. In experiment 3, we repeated the training procedure including all the augmentation techniques mentioned in Sec. IV.B (*RCSEHBG*) and obtained an F1-score of 0.613.

2) *Ensemble*: The goal of experiment 4 is to test whether the use of an ensemble of networks can improve the performance of the mitosis detector beyond the results obtained in experiment 3 with a single network. We trained and combined a set of CNNs, as explained in Sec. IV.C, and obtained an F1-score of 0.660. Furthermore, we analyzed its performance with respect to the detection threshold and observed a more robust behavior than that of the single CNN tested in experiment 3, illustrated in Fig. 7.

TABLE IV

INDEPENDENT EVALUATION OF THE PROPOSED METHOD PERFORMANCE IN THE THREE TASKS OF THE TUPAC CHALLENGE. COLUMNS TOP-1, TOP-2 AND TOP-3 CORRESPOND TO THE BEST PERFORMING SOLUTIONS IN THE PUBLIC LEADERBOARD, RESPECTIVELY

Dataset	Ground truth	Metric	Top-1	Top-2	Top-3	Proposed [95 c.i.]
TUPAC-test	Tumor grading	Kappa	0.567	0.534	0.462	0.471 [0.340, 0.603]
TUPAC-test	Proliferation score	Spearman	0.617	0.516	0.503	0.519 [0.477, 0.559]
TUPAC-aux-test	Mitosis location	F1-score	0.652	0.648	0.616	0.480

3) *Distillation*: The goal of experiments 5, 6 and 7 is to test whether knowledge distillation can effectively transfer the performance of the ensemble trained in experiment 4 to a single CNN. We trained three CNNs with γ set to 1.0, 0.8 and 0.6, respectively to experiments 5, 6 and 7. They all exhibited a similar performance to that of the ensemble (F1-score of 0.623, 0.628 and 0.636, respectively), with drastically less trainable parameters and superior performance to a single CNN trained without distillation (experiment 3). Experiment 7 resulted in *CNN5*, used in the following sections.

B. Comparison With the State of the Art

We evaluated the performance of our system in the three tasks of the TUPAC Challenge [9], and compared the results with those of top-performing teams, summarized in Tab. IV. We used *CNN5* with $\gamma = 0.6$ for all submissions, solely trained with the *TNBC-H&E* dataset. Notice that the authors do not have access to the ground truth data of *TUPAC-test* and *TUPAC-aux-test*. Our model predictions were independently evaluated by the organizers of TUPAC. This ensured fair and independent comparison with state of the art methods. For the first and second tasks, we tuned the hyper-parameters of the proposed whole-slide mitosis detector with the *TUPAC-train* dataset. We selected $\delta = 0.970$ to maximize the Spearman correlation between our mitotic count prediction and the ground truth proliferation score. Then, we tuned θ_1 and θ_2 to maximize the quadratic weighted Cohen’s kappa coefficient between our predicted tumor grade and the ground truth, obtaining $\theta_1 = 6$ and $\theta_2 = 20$. For the third task, we selected $\delta = 0.919$ to maximize the F1-score metric in the *TUPAC-aux-train* dataset. Spearman, kappa and F1-score are the evaluation metrics proposed in the TUPAC Challenge.

For the first task, we obtained a Kappa agreement of 0.471 with 95 confidence intervals [0.340, 0.603] between the ground truth tumor grading and our prediction on the *TUPAC-test* dataset. This performance is comparable to the top-3 entry in the leaderboard.

For the second task, we obtained a Spearman correlation of 0.519 with 95 confidence intervals [0.477, 0.559] between the ground truth genetic-based proliferation score and our prediction on the *TUPAC-test* dataset. This performance is comparable to the top-2 entry in the leaderboard.

For the third task, we obtained an F1-score of 0.480, with a precision and recall values of 0.467 and 0.494, respectively, by detecting individual mitotic figures in the *TUPAC-aux-test*. This performance is comparable to the top-7 entry in the leaderboard.

C. Observer Study: Precision of the Mitosis Detector

Due to the relatively low F1-score of 0.480 obtained in the *TUPAC-aux-test*, compared to the F1-score of 0.636 obtained in the *TUPAC-aux-train*, we investigated whether this difference could potentially be caused by a combination of inter-observer variability in the TUPAC reference standard, which was established by human observers, and lack of sufficient number of test samples. A resident pathologist manually classified the detections of *CNN5* on the *TUPAC-aux-test*, blinded to the patch labels. The observer indicated that 128 out of 181 detections contained mitotic figures, resulting in a precision of 0.707 for the detector. With this precision and assuming the recall suggested by the organizers, we would obtain an alternative F1-score of 0.581 in the *TUPAC-aux-test*. For the sake of completeness, the patches used in this experiment are depicted in the Appendix (Fig. 8).

VII. DISCUSSION AND CONCLUSION

To the best of our knowledge, this is the first time that the problem of noisy reference standards in training algorithms for mitosis detection in H&E WSIs was solved using immunohistochemistry. We validated our hypothesis that mitotic activity in PHH3 can be exploited to train a mitosis detector for H&E WSIs that is competitive with the state of the art. We proposed a method that combined (1) H&E-PHH3 restaining, (2) automatic image analysis and (3) registration to exhaustively annotate mitotic figures in entire H&E WSIs, for the first time. Only 2 hours of manual annotations per observer were needed to train the algorithm, delivering a dataset that was at least an order of magnitude larger than the publicly available one for mitosis detection. Using this method, the total number of annotated mitotic figures in H&E was solely limited by the number of restained slides available, not the amount of manual annotations. This is a very desirable property in the Computational Pathology field where manual annotations require plenty of resources and human expertise. Our work serves as a proof of concept to show that the combination of restaining, image analysis and registration can be used to automatically generate ground truth at scale when immunohistochemistry is the reference standard.

Staining variation between centers has long prohibited good generalization of algorithms to unseen data. In this work we applied a stain augmentation strategy that modifies the hematoxylin and eosin color channels directly, resulting in training samples with diverse and realistic H&E stain variations. Our experimental results indicate that the use of H&E-specific data augmentation and an ensemble of networks were key ingredients to drastically reduce the CNN’s generalization error to unseen stain variations, i.e. transferring the performance

obtained in WSIs from a single center to a cohort of WSIs from multiple centers. Furthermore, these results suggest that it is possible to train robust mitosis detectors without the need for assembling multicenter training cohorts or using stain standardization algorithms. More generally, we think that this combination of H&E-specific data augmentation and ensembling could benefit other applications where inference is performed on H&E WSIs, regardless of the tissue type.

High capacity CNNs typically exhibit top performance in a variety of tasks in the field of Computational Pathology, including mitosis detection. However, they come with a computational burden that can potentially compromise their applicability in daily practice. By using knowledge distillation, we massively reduced the computational requirements of the trained detector at inference time. In particular, we shrank the size of the distilled model 28 times (see Tab. III), with a negligible performance loss. This reduction combined with the fully convolutional design of the distilled network enabled us to perform efficient dense prediction at gigapixel-scale, processing entire TUPAC WSIs at $0.25\mu\text{m}/\text{pixel}$ resolution in 30-45 min.

On the task of individual mitosis detection in the *TUPAC-aux-test* set, we obtained different precision scores from the TUPAC organizers (0.467) and our observer (0.707). We attribute this disagreement to two factors: (1) the method used to annotate the images, and (2) the small number of samples in the test set. According to TUPAC organizers, two pathologists independently traversed the image tiles and identified mitotic figures. Coincident detections were marked as reference standard, and discording items were reviewed by a third pathologist. Notice that mitotic figures missed by both pathologists were never reported, potentially resulting in true mitotic figures not being annotated. This lack of recall could explain the high number of false positives initially detected by our network and later found to be true mitotic figures by an expert observer. Due to the small number of samples in the *TUPAC-aux-test* set (34 tiles), this effect can cause a large distortion in the F1-score. For reproducibility considerations, we have included all detections in the Appendix (Fig. 8). These results illustrate the difficulty of annotating mitotic figures manually, specifically in terms of recalling them throughout large tissue regions, and supports the idea of using PHH3 stained slides as an objective reference standard for the task of mitosis detection.

As a limitation of our method, we acknowledge the existence of some noise in the reference standard generated by analyzing the PHH3 WSIs and attribute it to three components: (1) the limited sensitivity of the PHH3 antibody (some late-stage mitotic figures were not highlighted, thus not even considered as candidates); (2) the limited specificity of the PHH3 antibody (many of the candidates turned out to be artifacts); and (3) the limited performance of CNN2 (F1-score of 0.9). This noise restricted our ability to detect small performance changes during training, potentially resulting in suboptimal model and/or hyper-parameter choices. More carefully PHH3 restaining process and improved training protocols could palliate this effect in the future.

In terms of future work, mitotic density at whole-slide level could be exploited to find the location of the tumor

hotspot, potentially resulting in significant speedups in daily practice. Additionally, the same metric could be used to study tumor heterogeneity, e.g. by analyzing the distribution of active tumor fronts within the sample. More generally, our work could be extended into other areas of Computational Pathology beyond mitosis detection in breast tissue by: (1) adopting the combination of slide restaining, automatic image analysis and registration to create large-scale training datasets where immunohistochemistry is the reference standard; and (2) validating the proposed stain augmentation strategy in other applications that analyze H&E WSIs.

APPENDIX

A. Theory of Color Representation

The Lambert-Beer law describes the relation between the amount of light absorbed by a specimen and the amount of stain present on it:

$$\frac{I_i}{I_{0,i}} = \exp(-Ac_i), \quad (1)$$

where I_i is the radiant flux emitted by the specimen, $I_{0,i}$ is the radiant flux received by the specimen, A is the amount of stain, c_i is the absorption factor, and subscript i indicates one of the RGB color channels.

Based on this law, we cannot establish a linear relationship between the relative absorption detected by a RGB camera ($I_i/I_{0,i}$) and the amount of stain present in a specimen. However, we can instead define the optical density (OD) per channel as follows:

$$OD_i = -\log \frac{I_i}{I_{0,i}} = -Ac_i. \quad (2)$$

Each OD vector describes a given stain in the OD -converted RGB color space. For example, measurements of a specimen stained with hematoxylin resulted in OD values of 0.18, 0.20 and 0.08 for each of the RGB channels, respectively [33].

By measuring the relative absorption for each RGB channel on slides stained with a single stain, Ruifrok and Johnson [33] quantified these OD vectors for hematoxylin, eosin and DAB (HED) stains. We can group these vectors into M , a 3 by 3 matrix representing a linear relationship between OD -converted pixels and the HED stain space. To achieve a correct balancing of the absorption factor for each stain, we divide each OD vector in M by its total length.

Therefore, a particular set of OD -converted pixels y can be described as follows:

$$y = xM, \quad (3)$$

where x is a 1 by 3 vector representing the amount of stain (e.g. hematoxylin, eosin and DAB) per pixel, and M is the normalized OD matrix for the given combination of stains. Since we are interested in obtaining x , we only need to invert M :

$$x = yM^{-1}. \quad (4)$$

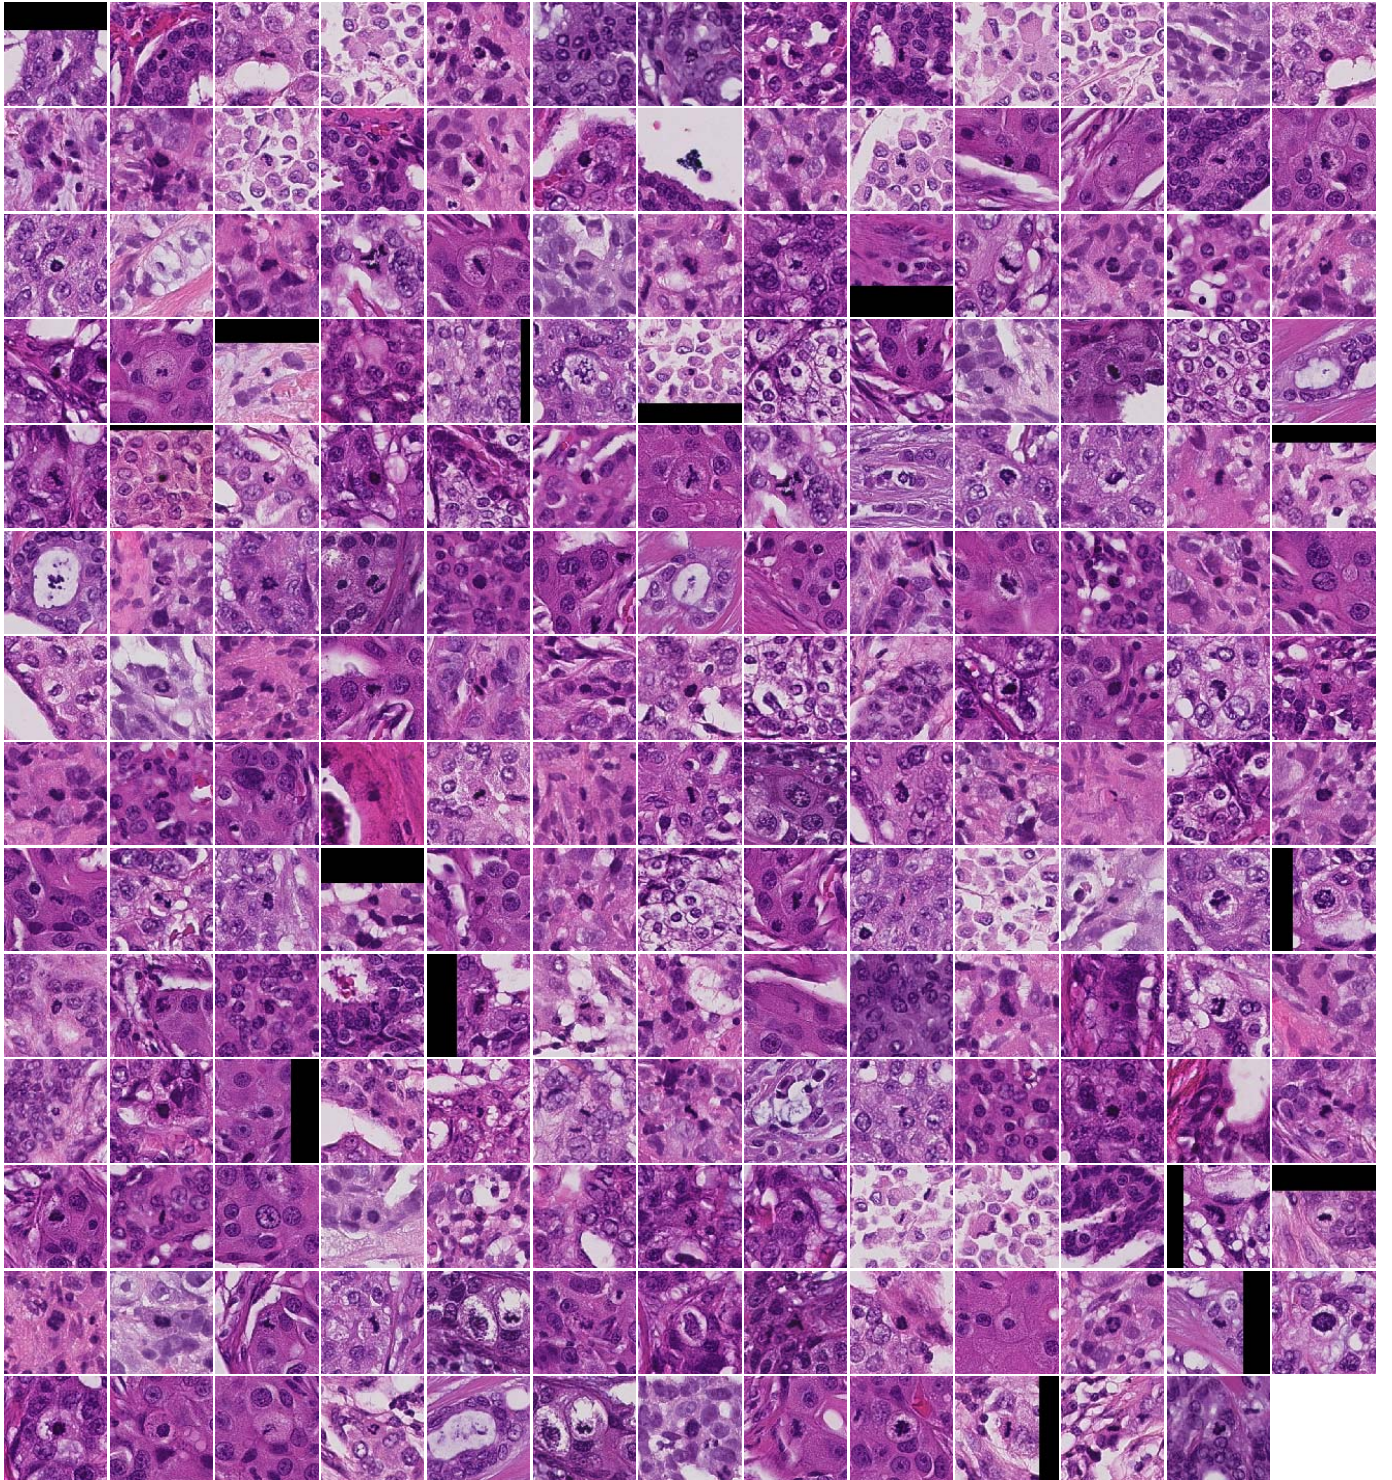


Fig. 8. Mitosis detections in the *TUPAC-aux-test* dataset identified by *CNN5*. These patches were classified as containing a mitotic figure or not by a resident pathologist. The observer classified 128 out of 181 detections as true positives, resulting in a precision score of 0.707 for the automatic detector.

B. Color Stain Augmentation Algorithm

Given an RGB image patch $P \in \mathbb{R}^{M \times M \times 3}$ reshaped into $P \in \mathbb{R}^{N \times 3}$ with N RGB pixels and the normalized *OD* matrix $M \in \mathbb{R}^{3 \times 3}$ for hematoxylin, eosin and DAB, we apply equations 2 and 4 to transform the patch from RGB to HED

color space as follows:

$$S = -\log(P + \epsilon)M^{-1}, \quad (5)$$

where $S \in \mathbb{R}^{N \times 3}$ is the transformed patch in HED color space and ϵ is a positive bias to avoid numerical errors. We simulate alternative stain intensities by stochastically mod-

ifying each stain component:

$$S'_i = \alpha_i S_i + \beta_i, \quad (6)$$

where $S' \in \mathbb{R}^{N \times 3}$ is the augmented patch in HED color space, subscript i represents each stain channel, α_i is drawn from a uniform distribution $U(1 - \sigma, 1 + \sigma)$, β_i is drawn from a uniform distribution $U(-\sigma, \sigma)$, and typically $\sigma = 0.05$.

To obtain an RGB representation of the augmented patch S' , we invert the operations described in equation 5:

$$P' = \exp(-S'M) - \epsilon, \quad (7)$$

where $P' \in \mathbb{R}^{N \times 3}$ is the augmented patch in RGB color space. Finally, we reshape P' into $P' \in \mathbb{R}^{M \times M \times 3}$ to match the original shape of the patch.

ACKNOWLEDGMENT

The authors would like to thank Mitko Veta, organizer of the TUPAC Challenge, for evaluating our predictions in the test set of the TUPAC dataset; the NVIDIA Corporation for donating a Titan X GPU card for our research; and the developers of Theano [35] and Lasagne [36], the open source tools that they used to run our deep learning experiments.

REFERENCES

- [1] H. J. G. Bloom and W. W. Richardson, "Histological grading and prognosis in breast cancer a study of 1409 cases of which 359 have been followed for 15 years," *Brit. J. Cancer*, vol. 11, no. 3, pp. 359–377, 1957.
- [2] C. W. Elston, I. O. Ellis, and S. E. Pinder, "Pathological prognostic factors in breast cancer," *Critical Rev. Oncol./Hematol.*, vol. 31, no. 3, pp. 209–223, 1999.
- [3] P. J. van Diest, E. van der Wall, and J. P. A. Baak, "Prognostic value of proliferation in invasive breast cancer: A review," *J. Clin. Pathol.*, vol. 57, no. 7, pp. 675–681, 2004.
- [4] I. Skaland, P. J. van Diest, E. A. M. Janssen, E. Gudlaugsson, and J. P. A. Baak, "Prognostic differences of World Health Organization–assessed mitotic activity index and mitotic impression by quick scanning in invasive ductal breast cancer patients younger than 55 years," *Hum. Pathol.*, vol. 39, no. 4, pp. 584–590, 2008.
- [5] S. Al-Janabi, H.-J. van Slooten, M. Visser, T. van der Ploeg, P. J. van Diest, and M. Jiwa, "Evaluation of mitotic activity index in breast cancer using whole slide digital images," *PLoS ONE*, vol. 8, no. 12, p. e82576, 2013.
- [6] L. Roux *et al.*, "Mitosis detection in breast cancer histological images an icpr 2012 contest," *J. Pathol. Inf.*, vol. 4, no. 1, p. 8, 2013.
- [7] M. Veta *et al.*, "Assessment of algorithms for mitosis detection in breast cancer histopathology images," *Med. Image Anal.*, vol. 20, no. 1, pp. 237–248, 2015.
- [8] L. Roux. (2014). *Mitosis Detection in Breast Cancer Histological Images*. [Online]. Available: <https://mitos-atypia-14.grand-challenge.org>
- [9] M. Veta. (2016). *Tumor Proliferation Assessment Challenge*. [Online]. Available: <http://tupac.tue-image.nl>
- [10] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Univ. Toronto, Toronto, ON, Canada, Tech. Rep., 2009. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/download;doi=10.1.1.222.9220&rep=rep1&type=pdf>.
- [11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2009, pp. 248–255.
- [12] D. C. Cireřan, A. Giusti, L. M. Gambardella, and J. Schmidhuber, "Mitosis detection in breast cancer histology images with deep neural networks," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervent.*, 2013, pp. 411–418.
- [13] M. Veta, P. J. van Diest, M. Jiwa, S. Al-Janabi, and J. P. W. Pluim, "Mitosis counting in breast cancer: Object-level interobserver agreement and comparison to an automatic method," *PLoS ONE*, vol. 11, no. 8, p. e0161286, 2016.
- [14] T. Ribalta, I. E. McCutcheon, K. D. Aldape, J. M. Bruner, and G. N. Fuller, "The mitosis-specific antibody anti-phosphohistone-H3 (PHH3) facilitates rapid reliable grading of meningiomas according to who 2000 criteria," *Amer. J. Surgical Pathol.*, vol. 28, no. 11, pp. 1532–1536, 2004.
- [15] C. M. Focke, K. Finsterbusch, T. Decker, and P. J. van Diest, "Performance of 4 immunohistochemical phosphohistone H3 antibodies for marking mitotic figures in breast cancer," *Appl. Immunohistochem. Mol. Morphol.*, vol. 26, no. 1, pp. 20–26, 2018.
- [16] I. Skaland *et al.*, "Validating the prognostic value of proliferation measured by phosphohistone H3 (PHH3) in invasive lymph node-negative breast cancer patients less than 71 years of age," *Breast Cancer Res. Treatment*, vol. 114, no. 1, pp. 39–45, 2009.
- [17] H. Colman *et al.*, "Assessment and prognostic significance of mitotic index using the mitosis marker phospho-histone H3 in low and intermediate-grade infiltrating astrocytomas," *Amer. J. Surgical Pathol.*, vol. 30, no. 5, pp. 657–664, 2006.
- [18] S. Fukushima *et al.*, "Sensitivity and usefulness of anti-phosphohistone-H3 antibody immunostaining for counting mitotic figures in meningioma cases," *Brain Tumor Pathol.*, vol. 26, no. 2, pp. 51–57, 2009.
- [19] K. Ikenberg, M. Pfaltz, C. Rakozy, and W. Kempf, "Immunohistochemical dual staining as an adjunct in assessment of mitotic activity in melanoma," *J. Cutaneous Pathol.*, vol. 39, no. 3, pp. 324–330, 2012.
- [20] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>
- [21] E. Zerhouni, D. Lányi, M. Viana, and M. Gabrani, "Wide residual networks for mitosis detection," in *Proc. IEEE 14th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2017, pp. 924–928.
- [22] K. Paeng, S. Hwang, S. Park, and M. Kim. (Dec. 2016). "A unified framework for tumor proliferation score prediction in breast histopathology." [Online]. Available: <https://arxiv.org/abs/1612.07180>
- [23] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361841517301135>
- [24] M. Macenko *et al.*, "A method for normalizing histology slides for quantitative analysis," in *Proc. IEEE Int. Symp. Biomed. Imag. Nano Macro (ISBI)*, Jun./Jul. 2009, pp. 1107–1110.
- [25] A. M. Khan, N. Rajpoot, D. Treanor, and D. Magee, "A nonlinear mapping approach to stain normalization in digital histopathology images using image-specific color deconvolution," *IEEE Trans. Biomed. Eng.*, vol. 61, no. 6, pp. 1729–1738, Jun. 2014.
- [26] B. E. Bejnordi *et al.*, "Stain specific standardization of whole-slide histopathological images," *IEEE Trans. Med. Imag.*, vol. 35, no. 2, pp. 404–415, Feb. 2016.
- [27] Y. Liu *et al.* (Mar. 2017). "Detecting cancer metastases on gigapixel pathology images." [Online]. Available: <https://arxiv.org/abs/1703.02442>
- [28] G. Hinton, O. Vinyals, and J. Dean. (Mar. 2015). "Distilling the knowledge in a neural network." [Online]. Available: <https://arxiv.org/abs/1503.02531>
- [29] J. N. Weinstein *et al.*, "The cancer genome atlas pan-cancer analysis project," *Nature Genet.*, vol. 45, no. 10, pp. 1113–1120, 2013.
- [30] T. O. Nielsen *et al.*, "A comparison of PAM50 intrinsic subtyping with immunohistochemistry and clinical prognostic factors in tamoxifen-treated estrogen receptor-positive breast cancer," *Clin. Cancer Res.*, vol. 16, no. 21, pp. 5222–5232, 2010.
- [31] P. Bándi *et al.* (Apr. 2017). "Comparison of different methods for tissue segmentation in histopathological whole-slide images." [Online]. Available: <https://arxiv.org/abs/1703.05990>
- [32] P. Y. Simard, D. Steinkraus, and J. C. Platt, "Best practices for convolutional neural networks applied to visual document analysis," in *Proc. ICDAR*, vol. 3, 2003, pp. 958–962.
- [33] A. C. Ruifrok and D. A. Johnston, "Quantification of histochemical staining by color deconvolution," *Anal. Quant. Cytol. Histol.*, vol. 23, no. 4, pp. 291–299, Aug. 2001.
- [34] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. (Dec. 2014). "Striving for simplicity: The all convolutional net." [Online]. Available: <https://arxiv.org/abs/1412.6806>
- [35] T. D. Team *et al.* (May 2016). "Theano: A Python framework for fast computation of mathematical expressions." [Online]. Available: <https://arxiv.org/abs/1605.02688>
- [36] S. Dieleman *et al.* (Aug. 2015). *Lasagne: First Release*. [Online]. Available: <http://dx.doi.org/10.5281/zenodo.27878>