MITOS-RCNN: Mitotic Figure Detection in Breast Cancer Histopathology Images Using Region Based Convolutional Neural Networks

Siddhant Rao

Abstract—Studies estimate that there will be 266,120 new cases of invasive breast cancer and 40,920 breast cancer induced deaths in the year of 2018 alone. Despite the pervasiveness of this affliction, the current process to obtain an accurate breast cancer prognosis is tedious and time consuming. It usually requires a trained pathologist to manually examine histopathological images and identify the features that characterize various cancer severity levels. We propose MITOS-RCNN: a region based convolutional neural network (RCNN) geared for small object detection to accurately grade one of the three factors that characterize tumor belligerence described by the Nottingham Grading System: mitotic count. Other computational approaches to mitotic figure counting and detection do not demonstrate ample recall or precision to be clinically viable. Our models outperformed all previous participants in the ICPR 2012 challenge, the AMIDA 2013 challenge and the MITOS-ATYPIA-14 challenge along with recently published works. Our model achieved an F- measure score of 0.955, a 6.11% improvement in accuracy from the most accurate of the previously proposed models.

Keywords—Object detection, histopathology, breast cancer, mitotic count, deep learning, computer vision.

I. INTRODUCTION

ONE in eight U.S. women will develop invasive breast cancer at some point in their lives, placing breast cancer as the second most commonly diagnosed form of cancer, regardless of gender [1]. The World Health Organization recommends the use of the Nottingham Grading System for tumor grading [2]. The Nottingham Grading System is derived from the assessment of three main morphological features: nuclear atypia, mitotic count and tubule formation. Nuclear atypia is described as the deformation of nuclei in a population of cells and is characterized by the following factors: size of nuclei, size of nucleoli, density of chromatin, thickness of nuclear membrane, regularity of nuclear contour, and anisonucleosis (size variation within a population of nuclei). Tubule formation is described as the percent of cancer cells that are in regular tubule formation. As the cancer becomes more belligerent, the tumor cells proliferate via mitosis (the process of cellular division), making the mitotic count of a tumor an important prognostic factor. For this study, we will be focusing on the most documented and salient feature involved in an accurate breast cancer prognosis: mitotic count. Mitotic count needs little to no professional interpretation, due to the simple metrics used to identify proliferation rates using the mitotic count per high power field (HPF's: the area visible under the maximum magnification power of a microscope) :

Siddhant Rao is with the Monta Vista High School, United States (e-mail: sanjaymk908@yahoo.com).

0-9 mitoses per 10 HPF's is low proliferation, 10-19 mitoses per 10 HPF's is moderate proliferation and more than 19 mitoses per 10 HPF's is severe proliferation.

Despite the prevalence of breast cancer, current methods for breast cancer prognosis are quite primitive. Trained pathologists are needed to examine hundreds of high power fields of histology images. Biopsies often take around two to ten days for results to return to the patient [1]. Given the growing number of breast cancer incidences [1], the traditional method for breast cancer prognosis is not sustainable. A computational approach would be a much more time and cost effective alternative, allowing for a streamlined breast cancer prognosis pipeline. This would allow for the deploying of pathological services to impoverished areas and the optimization of care centers globally.

(a) Prophase   (b) Metaphase

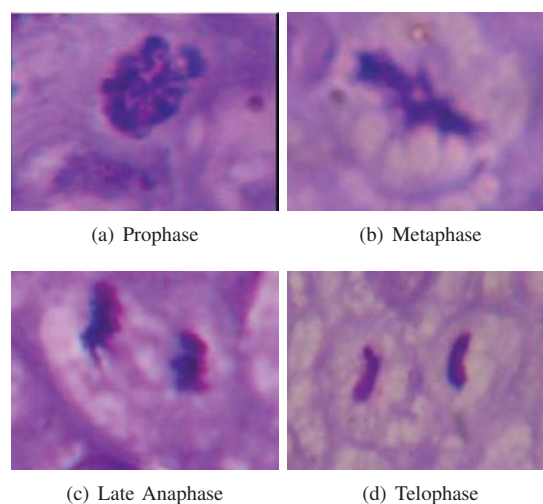(c) Late Anaphase   (d) Telophase

Fig. 1 Different stages of mitosis. The shape variation across stages can cause low mitotic figure detection accuracy

However, there are some complications that limit the accuracy of both computational and manual mitotic count extraction. Obtaining an accurate mitotic count is quite a challenge, as mitoses are often of low density, but high variation throughout HPF's [3]. Such variation is seen across the four phases of mitosis (prophase, metaphase, anaphase and telophase), with each phase having its own distinct size and shape (see Fig. 1). Mitotic figures in the anaphase or telophase stage of mitosis are often misclassified as 2 mitotic figures rather than 1. The low density of mitotic figures is evident in

the metrics used to classify various cancer severity levels using mitotic counts: 0-9 mitoses per 10 HPF's is low proliferation, 10-19 mitoses per 10 HPF's is moderate proliferation and more than 19 mitoses per 10 HPF's is severe proliferation. On average there are about 0-2 mitotic figures per HPF. Low density and high variation of mitotic figures makes scanning through hundreds of HPF's a tedious task when done manually and makes the practice susceptible to human error. For example, apoptotic cells (cells undergoing preprogrammed cell death) and other debris accumulated while preparing the tissue sample are often confused with mitoses due to their shaded, circular appearance. Irregularities in hematoxylin and eosin (H&E) staining across cancer research/treatment centers also add to the variation of breast cancer histopathology images.

Prior computational approaches to mitotic figure detection in breast cancer histopathological images within the scope of contests do not generalize well to new sets of data, resulting in relatively poor performance on an evaluation dataset. Outside the scope of participants in a mitotic figure detection challenge, several improvements have been made, but the methods are not accurate enough for clinical viability.

Deep learning is a growing field geared towards multi-scale pattern detection using deep neural network architectures. Adaptations of models like the convolutional neural network (CNN) can extract high level features from images to be used for object detection tasks like obtaining a mitotic count. One example of such a model is the Faster-RCNN proposed by [4], which uses features from an image to produce spatial coordinates for bounding boxes associated with certain classes.

We propose MITOS-RCNN: an adaptation of the Faster-RCNN model geared towards the automatic detection and counting of mitotic figures in breast cancer histopathology images. Our model was trained using the ICPR 2012, AMIDA 2013 and MITOS-ATYPIA-14 challenge datasets. We later compare the results of our models when fed sample images to those of previous works and demonstrate that our model significantly outperforms all other approaches.

II. Related Work

A. Deep Learning

Although deep learning methodologies have just recently begun to gain popularity, the underlying theory and applications have been present for quite some time. One of the earlier applications of deep learning for image analysis was the work done by [5] using CNN's for handwritten zip code classification. However, support vector machines outperformed CNN's during this time due to the lack of computational resources available for deep learning methodologies to be successful. Reference [6] improved upon the work done by [5] with the introduction of CNN's for general object image classification and outperformed all existing methods in the ImageNet Large Scale Visual Recognition Challenge, thus showing the promise of deep learning techniques. Reference [7] unified object classification tasks and object detection tasks with the original RCNN model. Models like the Fast-RCNN [8] and the Faster-RCNN [4] improved the speed at which

these models would be trained and evaluated, resulting in close to real-time object detection.

B. Mitotic Figure Detection

Prior contests have been held with the sole purpose of discovering novel approaches to detecting mitotic figures in histology images such as MITOS challenge at ICPR 2012 [9] and the AMIDA 2013 challenge [10]. The winners of MITOS contest, [11], utilized a deep, max-pooling CNN which operates on patches of pixels and their respective color channels and classifies those pixels as mitotic or not mitotic figures, ultimately achieving an F-measure score of 0.782. The model was trained to formulate features based on training images, contrasting the other contestants use of handcrafted features. Ciresan et al.'s work was one of the first applications of CNN's in a histopathological context.

The AMIDA 2013 challenge [10] proved to be quite similar, as contestants either employed classifiers (e.g. random forest classifiers) that relied upon hand-crafted features or utilized deep learning methodologies similar to those previously proposed in the MITOS challenge. Ciresan et al. [11] prevailed once again with the use of Multi Column Max-Pooling Convolutional Neural Networks (MCMPCNN). This new approach utilized a probabilistic representation of whether a pixel was a mitosis or not along with three 10-layer networks working in tandem. Reference [11] achieved an F-measure score of 0.611 with this approach.

Winners of the MITOS-ATYPIA-14 challenge achieved an F-measure score of 0.356 using a model called the Deep Cascaded Network [12] consisting of the following main steps: (1) candidate mitotic figure detection using a fully convolutional network and (2) discrimination classification of the detected mitotic figures candidates using a CNN.

Recent advances in mitotic figure detection outside the scope of contests have been made by various works. For example, [13] employed a CNN paired with a crowd-sourced learning architecture and achieved an F-measure of 0.433. Additionally, Saha et al. [14] use both hand-crafted and learned features in their proposed model, achieving an impressive F-measure score of 0.900.

Our proposed approach builds upon the most successful prior works by utilizing a modified Faster-RCNN tuned for the detection of small objects, which matches the speed of previous CNN implementations but with more accurate detections.

III. Materials

A. Dataset Description

Our evaluation set of data consists of 100 samples of each of the "mitotic figure" and "not mitotic figure" classes spread across 187 HPF subsections.

IV. Methods

A. Data Preprocessing and Augmentation

Our model takes in only 299x299 px images as input and resizes any input to this dimension. We split each

World Academy of Science, Engineering and Technology
International Journal of Medical and Health Sciences
Vol:12, No:10, 2018

TABLE I
THE DIMENSIONS, RESOLUTION AND NUMBER OF FRAMES FROM EACH OF THE TWO SCANNERS: THE APERIO SCANSCOPE XT AND THE HAMAMATSU NANOZOOMER 2.0-HT

Scanner	Dimensions of x40 frame (px)	Res. at x40 (µm/px)	# of Frames
Aperio	1539 * 1376	0.2455	2622
Hamamatsu	1663 * 1485	0.2273	2016

image into 16 equal subsections to make sure minimal downsampling of each image takes place. This allows for more "attention" to small-scale features. Since the Faster-RCNN outputs bounding box spatial coordinates, we introduced a new dataset, MITOS-BOXES, using the preexisting centroid coordinates and annotating box regions by hand for each mitotic and non-mitotic figure. Due to the relatively small size of our dataset, we introduced artificially augmented versions of existing data to create more samples for the model to learn from and to prevent overfitting. We rotated all images by preset values (90°, 180°, 270°), as we believed that keeping raw pixel data intact would be beneficial for the model. Due to the inconsistency of staining techniques across the data sources, all image data was normalized via a procedure described by [15] (original samples were kept in dataset). The final size of the dataset was 37,104 HPF's.

B. Proposed Model

Our proposed model (see Fig. 2 for diagram of architecture) is a variation of [4] proposed Faster-RCNN model and is composed of two main modules: (1) a region proposal network (RPN) which returns regions of interest (ROI's) and (2) a detection network which classifies and discriminates regions of interests while performing a bounding box regression, ultimately returning spatial coordinates with associated classes. We utilized the VGG-16 model [16] as our base feed-forward CNN in order to extract powerful hierarchical features from an input image. We refer to the 'conv1_3', 'conv2_3', 'conv3_3', 'conv4_3' and 'conv5_3' layers of the VGG-16 network as $conv_1$, $conv_2$, $conv_3$, $conv_4$, $conv_5$, respectively. Region proposals are generated during the two-stage top-down cascade multi-scale proposal generation process, where features from both the $conv_3$ and $conv_4$ feature maps are aggregated and used by two separate RPN's: RPN_1 and RPN_2. The generated region proposals are then fed into two sibling fully connected layers (FC_{reg} and FC_{cls}) to regress bounding box spatial coordinates and classify the generated region proposals as being either "mitotic figure" or "not mitotic figure".

1) Small Scale Object Detection:
[17] proved that with the standard Faster-RCNN architecture, the minimum detectable object size was around 44px. This is due to the loss of information as the feature map representation dimensionality is reduced via the pooling layers within the feed-forward network. Essentially, the later $conv$ layer feature maps contain rich abstract-level features which are far too coarse to be used for extraction of features pertaining to small objects. A 44px object detection threshold is not optimal for the detection of mitoses, as the average size of the mitotic figures in our MITOS-BOXES dataset was around 30px. To avoid this issue we omit the $conv_5$ layer and utilize the feature maps from

only $conv_3$ and $conv_4$, allowing for the minimum detectable object size to be 15px and 22px, respectively [17].

2) Two-Stage Top-Down Cascade Multi-Scale Proposal Generation:
Our adapted RCNN model generates multi-scale proposals at two distinct stages of the network using features aggregated from multiple convolutional layers in a top-down manner in order to obtain more refined proposals for specifically small-scale objects. This allows the model to use semantic knowledge from both higher and lower level features to produce accurate object detections. RPN_1 follows the $conv_4$ layer and generates around 15k proposals (most of which are discriminated via NMS thresholding of 0.7). RPN_2 utilizes features aggregated from a feature map consisting of both $conv_3$ and $conv_4$ concatenated. To obtain a concatenated feature map, the $conv_4$ feature map is upsampled by a subsequent deconvolutional layer in order to obtain a resolution matching that of $conv_3$. Then, we normalize each layer using L2 norm, concatenate the two feature maps to obtain $conv_{3+4}$ and reduce the $conv_{3+4}$ feature map to a dimension of 256x1x1. RPN_2 utilizes inputs from two sources: (1) the proposals from RPN_1 and (2) the output of a "sliding window" operating on the $conv_{3+4}$ feature map with a scale of 64^2px and a 1:1 aspect ratio. Proposals are further refined by the model by discriminating low-quality proposals generated by RPN_2 via NMS thresholding of 0.7 and then fed into an ROI pooling layer to normalize region proposal scales.

3) Detection Network:
The detection network of our model is identical to that of the RCNN's proposed by [4] and [8]. Two sibling output layers make up this detection network: a bounding box classification layer and a bounding box regression layer. Both networks are fully-connected layers and receive the refined region proposals from the RPN_2 layer of the model after the proposal scales have been normalized by the ROI pooling layer. The bounding box classification layer, FC_{cls}, outputs a softmax probability distribution, $p = (p_0...p_k)$, over the 0th ...kth classes for every region proposal, where the softmax function is defined as:

$$P(c_i|x) = \frac{\exp(y_i)}{\sum_j^N \exp(y_j)}. \tag{1}$$

Given an input, x, the softmax function computes the probability of a class, c_i, using the classification score for the ith class, y_i, and the classification scores for all the classes.

The fully connected layer used for bounding box regression, FC_{reg}, outputs a vector of regression offsets, $t^k = (t_x^k, t_y^k, t_w^k, t_h^k)$, specifying a scale-invariant transformation (of the top-left corner of the box (x, y) and the width and height of the box (w, h)) to the input region proposal coordinates for each of the k classes.

World Academy of Science, Engineering and Technology
International Journal of Medical and Health Sciences
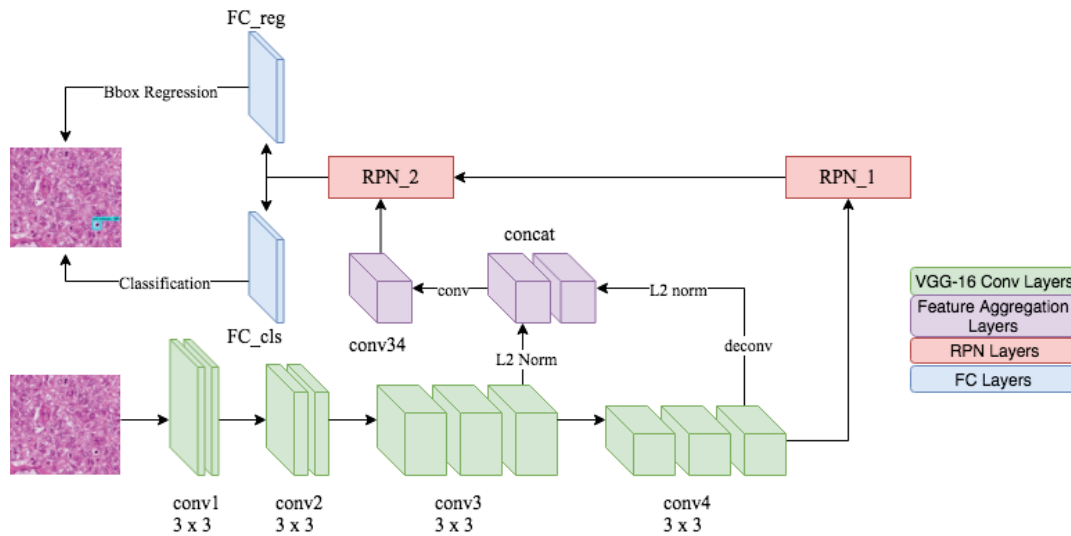Vol:12, No:10, 2018

Fig. 2 The architecture of our proposed MITOS-RCNN model, from the VGG-16 convolutional layers to the fully-connected layers outputting the final detections

4) Training: The loss function being minimized during the training process takes into account the loss from both modules of the detection network - the regression network, L_{reg}, and the classification network, L_{cls} - and is defined as:

$$L(p_i, t_i) = \frac{1}{N_{cls}} \sum_i L_{cls}(p_i, p_i^*) + \lambda \frac{1}{N_{reg}} \sum_i p_i^* L_{reg}(t_i, t_i^*),$$
(2)

where i is the anchor index (anchors are synonymous with bounding box proposals), p_i is the "objectness" of the anchor i, p_i^* is the predicted "objectness" of the anchor i, t_i is the coordinate vector of the bounding box prediction, t_i^* is the coordinate vector for the ground truth bounding box with a positive anchor, N_{cls} is the batch size, N_{reg} is the number of anchors, L_{cls} is the log loss over the object and not object classes, $L_{reg} = R(t_i - t_i^*)$ (R is the robust loss function described by [8]) and λ is a balancing hyperparameter (set to 10 in our implementation).

Our model was trained in the same fashion as the Faster-RCNN [4]. Weights in all layers of the RCNN were initialized from a zero-mean Gaussian distribution with standard deviation of 0.01, while all VGG-16 layers remained with their pre-trained weights. We trained the fine-tuned the RCNN model with our proposed MITOS-BOXES dataset using standard stochastic gradient descent (SGD) [5]. Batch size was kept to 10 images. The model was trained for 60,000 mini-batches with a learning rate of 0.001 and then 20,000 mini-batches with a learning rate of 0.0001. A momentum of 0.9 and a weight decay of 0.0005 was used.

C. Implementation Details

Both the training and testing process were performed on the Google Cloud Platform ML Engine. We used 5 NVIDIA Tesla K80 GPUs and 3 parameter servers in order to distribute the training process. The testing process only required a single NVIDIA Tesla K80 GPU. Dataset augmentation/preprocessing and model evaluation was done locally on a computer running

TABLE II
COMPARATIVE PERFORMANCE OF OUR PROPOSED RCNN ADAPTATION VERSUS THE STANDARD FASTER-RCNN MODEL AND A FASTER-RCNN MODEL WITH ONLY $conv_4$ FEATURES

Method	F_1 score
Faster-RCNN [4]	0.502
Faster-RCNN w/ $conv_4$	0.709
Proposed	**0.955**

the MacOS High Sierra operating system with a 2.7 GHz Intel Core i5 processor and 8 GB of RAM. Our implementation was in the Tensorflow machine learning framework [18].

V. RESULTS

See Fig. 3 for sample detections.

A. Metrics

We used the F-measure (or F_1 score) score as the benchmark metric for our model. The F-measure score is defined as the harmonic mean of precision and recall:

$$F_1 = 2 * \frac{precision * recall}{precision + recall},$$
(3)

where precision and recall are defined as:

$$precision = \frac{TP}{TP + FP},$$
(4)

$$recall = \frac{TP}{TP + FN},$$
(5)

and TP is the number of true positives, FP is the number of false positives and FN is the number of false negatives. Reference [3] describe a true positive detection as a detection that is at most $8\,\mu m$ from the centroid of a ground truth mitosis.
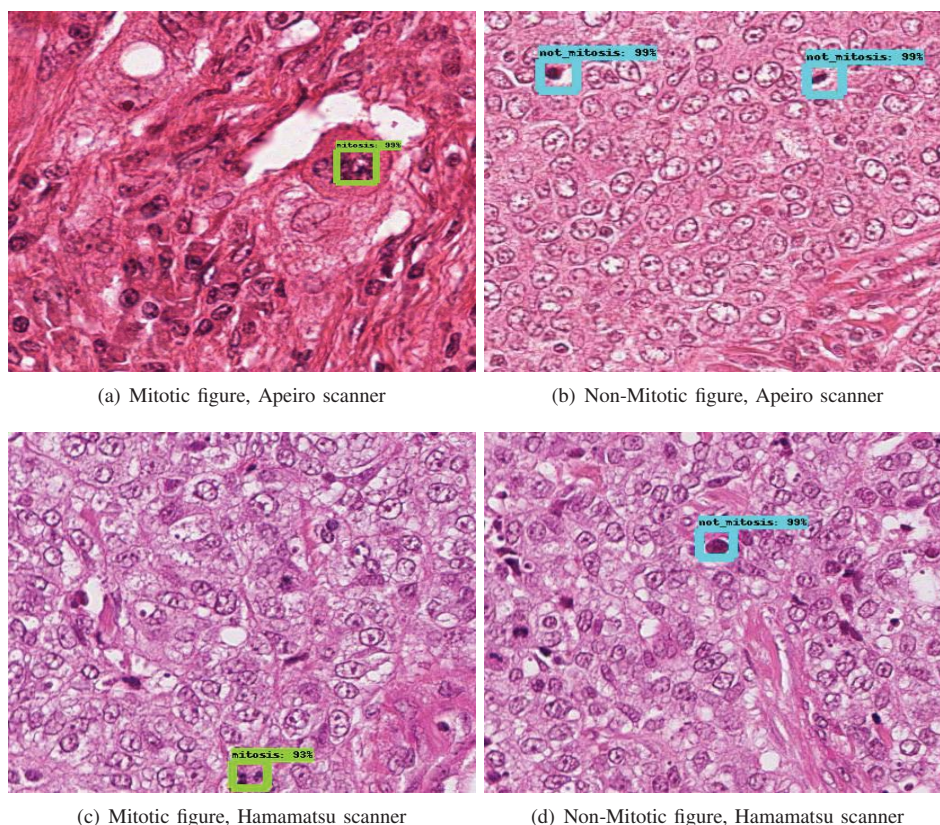
World Academy of Science, Engineering and Technology
International Journal of Medical and Health Sciences
Vol:12, No:10, 2018

(a) Mitotic figure, Apeiro scanner

(b) Non-Mitotic figure, Apeiro scanner

(c) Mitotic figure, Hamamatsu scanner

(d) Non-Mitotic figure, Hamamatsu scanner

Fig. 3 Sample detections of "mitotic figures" and "not mitotic figures" on HPF subsections from both an Apeiro scanner and a Hamamatsu scanner

| | | Ground Truth | | |
		Positive	Negative	Total
Prediction	Positive	72	31	103
	Negative	28	69	97
	Total	100	100	200
		Ground Truth		
		Positive	Negative	Total
Prediction	Positive	53	58	111
	Negative	47	42	89
	Total	100	100	200
		Ground Truth		
		Positive	Negative	Total
Prediction	Positive	96	5	101
	Negative	4	95	99
	Total	100	100	200

Fig. 4 Confusion matrices (from top to bottom): the standard Faster-RCNN, the Faster-RCNN with only $conv_4$ features, and our proposed model

TABLE III
COMPARATIVE RESULTS INVOLVING RECENTLY PUBLISHED WORKS AND CONTEST WINNERS

Reference	F_1 score
[11]	0.782
[10]	0.611
[12]	0.356
[13]	0.433
[14]	0.900
Proposed	**0.955**

method, we implemented a standard Faster-RCNN along with a Faster-RCNN only using $conv_4$ features and trained the models in a fashion identical to that of [4]. Table II shows the performance of our proposed model alongside the other 2 benchmark models. Our custom region proposal network improved upon the standard Faster-RCNN model by 90% and the Faster-RCNN with only $conv_4$ features by 35%.

2) Comparison with Previously Proposed Methods: Table III shows the results of our model in comparison to recently published works and contest winners. Our approach was 6.22% more accurate than the previous high F_1 score of 0.900 achieved by the model proposed by [14]. Our method outperformed all previously proposed approaches which utilized both handcrafted features and deep learning methodologies.

3) Computation Time: Most previous works do not document the time it takes for model evaluation. We found that, on average, our model took 0.5 seconds to process 1 HPF. Reference [14] report that their model took 0.3 seconds per HPF. This increase in our model's forward propagation time can be attributed to the increased complexity of our model due to the need for 2 RPN's.

B. Comparative Results

1) Two-Stage Top-Down Cascade Multi-Scale Proposal Generation Results: To display the efficacy of our custom two-stage top-down cascade multi-scale proposal generation

VI. DISCUSSION

A. Significance

Computerized extraction of mitotic counts from HPF's allows for a streamlined breast cancer prognosis pipeline. A

World Academy of Science, Engineering and Technology
International Journal of Medical and Health Sciences
Vol:12, No:10, 2018

biopsy taking 2 - 10 days [1] can be reduced significantly if manual mitotic figure counting was omitted from the current prognostic pipeline. Our contribution advances the state-of-the-art in computerized breast cancer prognosis, hopefully towards fully automated breast cancer prognosis in clinical practice.

Outside the scope of computerized medical imaging, the detection of small-scale objects is a practice applicable to many problems. Reference [17] proposed that his model be used for company logo detection in images where logos make up small fractions of the image. Satellite images also have many small artifacts or landmarks scattered around large, high-resolution images. Detecting such small objects with great accuracy is made possible by our proposed model. The future applications are abundant and promising.

B. Limitations

While models utilizing learned features may be more accurate than their counterparts relying on hand-crafted features, deep learning models of this scale have their limitations due to their computationally-exhaustive nature. Training a deep learning model like our proposed model requires GPU's in order to rapidly calculate extensive amounts of large-scale matrix operations along with large sets of varied, yet consistent data. Without multiple GPU's the time to complete training iterations can increase exponentially to infeasible values. Due to the advent of cloud computing power for deep learning models, we were able to address the issue of dealing with computationally taxing operations. However, there were more complications regarding the data which the model was trained, tested and evaluated on. While our dataset sources utilized the same set of tissue scanners (meaning relatively similar image characteristics), the histopathological images were from different labs using different staining protocols. Although we normalized the stain color across all images in our dataset, there would still be stain irregularities across our dataset which could result in increased false-positive or false-negative detections. Since our model was trained on data originally annotated by trained pathologists, our model was subject to human biases and error. The same variation of mitotic figures described in Section I also affects pathologists, meaning that misclassification are quite likely. Introducing a "not mitotic figure" class to account for low-confidence pathologically-annotated mitotic figures was an attempt to solve this problem, but there is no guarantee that pathologists correctly identified all mitotic figures with relatively high confidence. Apart from the inconsistencies or irregularities within our dataset, the size of our dataset is a pertinent problem for deep learning models. Many large-scale models rely on massive datasets during the training process. We artificially increased the dataset size and variation through augmentation techniques, but our dataset was still quite small compared to the size of the datasets utilized by other works. For example, the VGG-16 model (our base feed-forward network) was trained on a dataset with upwards of a million samples of image data [16].

VII. Conclusion

We propose a variant of the Faster-RCNN architecture in order to detect mitotic figures in breast cancer histopathological images with great speed and accuracy. Our two-stage top down multi-scale region proposal generation process enables our model to detect small objects such as the mitotic figures. Our results reinforce the strength of our proposed model in comparison to previously proposed works. Our proposed model achieved an F-measure score of 0.955, the highest accuracy achieved to date. Our model utilizes purely learned features to detect mitotic figures, thus displaying the strength of learned features compared to the hand-crafted features used by other works.

References

[1] M. Ghoncheh, Z. Pournamdar, and H. Salehiniya, "Incidence and mortality and epidemiology of breast cancer in the world," vol. 17, pp. 43–46, 06 2016.

[2] C. W Elston and I. Ellis, "Pathological prognostic factors in breast cancer. i. the value of histological grade in breast cancer: experience from a large study with long-term follow-up. c. w. elston & i. o. ellis. histopathology 1991; 19; 403-410. author commentary," vol. 41, pp. 151–2, discussion 152, 10 2002.

[3] L. Roux and D. Racoceanu, "Mitos & atypia detection of mitosis and evaluation of nuclear atypia score in breast cancer histological images," 06 2014.

[4] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in Advances in Neural Information Processing Systems 28, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Curran Associates, Inc., 2015, pp. 91–99. (Online). Available: http://papers.nips.cc/paper/5638-faster-r-cnn-towards-real-time-object-detection-with-region-\proposal-networks.pdf.

[5] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," Neural Computation, vol. 1, no. 4, pp. 541–551, Dec 1989.

[6] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1, ser. NIPS'12. USA: Curran Associates Inc., 2012, pp. 1097–1105. (Online). Available: http://dl.acm.org/citation.cfm?id=2999134.2999257.

[7] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," ArXiv e-prints, nov 2013.

[8] R. B. Girshick, "Fast R-CNN," CoRR, vol. abs/1504.08083, 2015. (Online). Available: http://arxiv.org/abs/1504.08083.

[9] L. Roux, D. Racoceanu, N. Lomenie, M. Kulikova, H. Irshad, J. Klossa, F. Capron, C. Genestie, G. Le Naour, and M. N Gurcan, "Mitosis detection in breast cancer histological images an icpr 2012 contest," vol. 4, p. 8, 05 2013.

[10] M. Veta, P. van Diest, S. Willems, H. Wang, A. Madabhushi, A. Cruz-Roa, F. Gonzalez, A. Larsen, J. Vestergaard, A. Dahl, D. Cirean, J. Schmidhuber, A. Giusti, L. Gambardella, F. Tek, T. Walter, C. Wang, S. Kondo, B. Matuszewski, F. Precioso, V. Snell, J. Kittler, T. de Campos, A. Khan, N. Rajpoot, E. Arkoumani, M. Lacle, M. Viergever, and J. Pluim, "Assessment of algorithms for mitosis detection in breast cancer histopathology images," Medical Image Analysis, vol. 20, no. 1, pp. 237–248, 2 2015.

[11] D. C. Cireşan, A. Giusti, L. M. Gambardella, and J. Schmidhuber, "Mitosis detection in breast cancer histology images with deep neural networks," in Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013, K. Mori, I. Sakuma, Y. Sato, C. Barillot, and N. Navab, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 411–418.

[12] H. Chen, Q. Dou, X. Wang, J. Qin, and P. Heng, "Mitosis detection in breast cancer histology images via deep cascaded networks," 2016. (Online). Available: https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11788/11717.

World Academy of Science, Engineering and Technology
International Journal of Medical and Health Sciences
Vol:12, No:10, 2018

[13] S. Albarqouni, C. Baur, F. Achilles, V. Belagiannis, S. Demirci, and N. Navab, "Aggnet: Deep learning from crowds for mitosis detection in breast cancer histology images," *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1313–1321, May 2016.

[14] M. Saha, C. Chakraborty, and D. Racoceanu, "Efficient deep learning model for mitosis detection using breast histopathology images," vol. 64, 12 2017.

[15] M. Macenko, M. Niethammer, J. Marron, D. Borland, J. Woosley, X. Guan, C. Schmitt, and N. Thomas, "A method for normalizing histology slides for quantitative analysis," in *Proceedings - 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, ISBI 2009*, 2009, pp. 1107–1110.

[16] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 09 2014.

[17] C. Eggert, S. Brehm, A. Winschel, D. Zecha, and R. Lienhart, "A closer look: Small object detection in faster r-cnn," in *2017 IEEE International Conference on Multimedia and Expo (ICME)*, vol. 00, July 2017, pp. 421–426. (Online). Available: doi.ieeecomputersociety.org/10.1109/ICME.2017.8019550.

[18] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, "Tensorflow: A system for large-scale machine learning," in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 2016, pp. 265–283. (Online). Available: https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf.