

Detecting Mitoses with a Convolutional Neural Network for MIDOG 2022 Challenge

Hongyan Gu¹[0000-0001-8962-9152], Mohammad Haeri²[0000-0001-6055-9779],
Shuo Ni¹[0000-0002-4125-2373], Christopher Kazu Williams³[0000-0003-3253-7630],
Neda Zarrin-Khameh⁴[0000-0001-7506-156X], Shino
Magaki³[0000-0003-0433-5759], and Xiang ‘Anthony’ Chen¹[0000-0002-8527-1744]

¹ University of California, Los Angeles, USA

² University of Kansas Medical Center, USA

³ UCLA David Geffen School of Medicine, USA

⁴ Baylor College of Medicine, USA

xac@ucla.edu

Abstract. This work presents a mitosis detection method with only one vanilla Convolutional Neural Network (CNN). Our method consists of two steps: given an image, we first apply a CNN using a sliding window technique to extract patches that have mitoses; we then calculate each extracted patch’s class activation map to obtain the mitosis’s precise location. To increase the model performance on high-domain-variance pathology images, we train the CNN with a data augmentation pipeline, a noise-tolerant loss that copes with unlabeled images, and a multi-rounded active learning strategy. In the MIDOG 2022 challenge, our approach, with an EfficientNet-b3 CNN model, achieved an overall F1 score of 0.7323 in the preliminary test phase, and 0.6847 in the final test phase (task 1). Our approach sheds light on the broader applicability of class activation maps for object detections in pathology images.

Keywords: Mitosis detection · Domain shift · Convolutional neural network · Class activation map.

1 Introduction

Mitotic activity is a crucial histopathological indicator related to cancer malignancy and patients’ prognosis [7]. Because of its importance, a considerable amount of literature has proposed datasets [2,3] and deep learning models [11,12] for mitosis detection. To date, a number of state-of-the-art deep learning approaches detect mitoses with the segmentation task [10,16]. However, these methods usually require generating pixel-level segmentation maps as ground truth. Another school of techniques utilizes object detection with a two-stage setup — a localization model (*e.g.*, RetinaNet) is employed for extracting interest locations, followed by a classification model to tell whether these locations have mitoses [2,11,12]. Such a two-stage setup was reported to improve the performance of mitosis detection compared to that with the localization model only [2].

Since adding a classification model can improve the performance, we argue that using only one CNN model for mitosis detection is also viable. Because CNNs cannot directly report the location of mitosis, previous works either modified the structure of CNNs [6], or used CNNs with a small input size to reduce the localization errors [14]. Instead, our approach extracts the location of mitoses with the class activation map (CAM) [17], which allows CNNs to accept a larger input size for more efficient training. Also, our approach is model-agnostic, and can work with vanilla CNNs because calculating CAMs does not require changing the network structure.

We validated our proposed method in Mitosis DOmain Generalization (MIDOG) 2022 Challenge [1]. The challenge training set consists of 403 Hematoxylin & Eosin (H&E) regions of interest (ROIs, average size= 5143×6860 pixels), covering six tumor types scanned from multiple scanners. 354/403 ROIs have been labeled and have 9,501 mitotic figures. The preliminary test set includes 20 cases from four tumor types, and the final test set has 100 independent tumor cases from ten tumor types. Given the dataset’s high variance, we employed three techniques to improve the CNN’s performance:

1. An augmentation pipeline with balanced-mixup [8] and stain augmentation [14];
2. The utilization of unlabeled images for training (treated as negative), and training with the Online Uncertainty Sample Mining (OUSM) [15] to gain robustness with noisy labels;
3. A multi-rounded, active learning training strategy that adds false-positive, false-negative, and hard-negative patches after each round of training.

2 Methods

2.1 Extracting Patches for Initial Training

We randomly used $\sim 90\%$ of the image instances in the MIDOG 2022 Challenge to generate the training set and $\sim 10\%$ for the validation set. To maximally utilize the dataset, we included unlabelled images in the training set and treated them as negative images (*i.e.*, no mitoses inside). For each image, we extracted patches with the size of $240 \times 240 \times 3$ pixels surrounding the location of each annotation (provided by the challenge) and placed them into the train/validation set.

2.2 Model Training

We trained an EfficientNet-b3 [13] model (input size: $240 \times 240 \times 3$) with pre-trained ImageNet weights. Here, we tried to improve model performance by constructing an online data augmentation pipeline. The pipeline includes general image augmentation techniques, including random rotation, flip, elastic transform, grid distortion, affine, color jitter, Gaussian blur, and Gaussian noise. Besides, we added two augmentation methods – stain augmentation [14] and




Fig. 1. (a) Illustration of the active learning training strategy used in this work; (b) Examples of augmented patches according to our augmentation pipeline; (c) Overall data processing pipeline our approach: detecting mitosis using a convolution neural network and the class activation map.

balance-mixup [8] – to deal with the domain shift in pathology images. Examples of augmented patches are shown in Figure 1(b). The model was trained with an SGD optimizer with momentum 0.9, a Cosine Annealing learning rate scheduler with warm restart ($\text{max LR}=6 \times 10^{-4}$). Since we treated all unlabeled images as negative, we further used an OUSM [15] + COnsistent RAnk Logits (CORAL) loss [4] to deal with noisy labels. Each round of training had 100 epochs, and we selected the model with the highest F1 score on the validation set for inferencing.

2.3 Inferencing

We slid the trained EfficientNet on train and validation images with window size 240×240 and step-size 30. We then cross-referenced the CNN predictions with the ground truth. Here, we define a positive window classification as a true-positive if mitoses were inside the window and a false-positive otherwise. We further define false-negative if no positive windows surround a mitosis annotation.

2.4 Incrementing the Patch Dataset with Active Learning

We employed a multi-round active learning process to boost the performance of the EfficientNet (Figure 1(a)). Each round starts with the model training on the current train/validation set (Section 2.2). Then, the best model is selected and applied to the images (Section 2.3). After that, false-positive, false-negative, and hard-negative patches are added to the train/validation set. The procedure

was repeated six times until the model’s F1 score on the validation set does not increase. Eventually, there are 103,816 patches in the final training set, and 23,638 in the validation set.

2.5 Extracting Mitosis Locations with CAMs

We used the best model from the final round in Section 2.4 for the test images. A window with a CNN probability > 0.84 was considered positive, and non-maximum suppression with a threshold of 0.22 was used to mitigate the overlapping windows. For each positive window, we calculated the CAM with GradCAM++ [5], and extracted the hotspot’s centroid as the mitosis location (Figure 1(c)). Examples of CAMs are shown in Figure 2(a). Specific numbers and thresholds were selected according to the best F1 performance from the validation images.

3 Results

On the preliminary test phase of the MIDOG 2022 Challenge, our approach achieved an overall F1 score of 0.7323, with 0.7313 precision and 0.7333 recall. In the final test phase of the challenge (task 1), our approach achieved the F1 score of 0.6847 (precision:0.7559, recall: 0.6258). In sum, our approach is 2.34% higher than the baseline RetinaNet approach regarding the overall F1 score in the preliminary test phase, but 4.21% lower in the final test phase. Please refer to the Grand-Challenge Leader-board⁵ for more details of the test result.




Fig. 2. Examples of positive patches detected by CNN and corresponding CAMs generated by GradCAM++ from image 240.tiff. (a) Patch samples and CAMs if there is one mitosis inside (pointed by the green arrow); (b) Patch samples and CAMs if there are multiple mitoses inside. Note that generated CAMs might not report strong signals for some mitoses.

⁵ <https://midog2022.grand-challenge.org/evaluation/final-test-phase-task-1-without-additional-data/leaderboard/>

4 Discussion & Conclusion

Although CAMs are primarily used for explaining CNN classifications, we demonstrate their potential usage to help detect mitoses in H&E images. We believe our approach has the potential to work with mobile/edge devices, which have limited computational power and are not optimized for special-structured deep-learning models.

It is noteworthy that CAMs might fail to highlight all mitoses when there are multiple in an image. As shown in Figure 2(b), if there were multiple mitoses inside one positive patch, CAMs might report a strong signal for one mitosis, while giving weak signals for others. However, we argue that our approach *as-is* can still be helpful in diseases where mitosis is not a high-prevalent histological pattern (*e.g.*, meningiomas [9]).

On the other hand, we believe that the limitation is partly caused by using classification models for localization tasks: given a patch with multiple mitoses, a CNN only needs to find at least one to predict it as positive. As such, other mitoses are more or less ignored. To compensate for the limitation, we suggest future works to improve the quality of CAMs by aligning them with ground-truth mitosis location maps. For example, a “CAM-loss” might be designed to penalize the misalignment between the CAM and mitosis location heatmaps [18]. Different from the segmentation task, generating mitosis location maps does not require pixel-level segmentation masks: simply applying a Gaussian kernel over the locations of mitoses can generate a mitosis location map similar to Figure 2.

5 Acknowledgement

This work was funded in part by the Office of Naval Research under grant N000142212188.

References

1. Aubreville, M., Bertram, C., Breininger, K., Jabari, S., Stathonikos, N., Veta, M.: Mitosis domain generalization challenge 2022. In: 25th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2022) (2022). <https://doi.org/10.5281/zenodo.6362337>
2. Aubreville, M., Bertram, C.A., Donovan, T.A., Marzahl, C., Maier, A., Klopffleisch, R.: A completely annotated whole slide image dataset of canine breast cancer to aid human breast cancer research. *Scientific data* **7**(1), 1–10 (2020)
3. Bertram, C.A., Aubreville, M., Marzahl, C., Maier, A., Klopffleisch, R.: A large-scale dataset for mitotic figure assessment on whole slide images of canine cutaneous mast cell tumor. *Scientific data* **6**(1), 1–9 (2019)
4. Cao, W., Mirjalili, V., Raschka, S.: Rank consistent ordinal regression for neural networks with application to age estimation. *Pattern Recognition Letters* **140**, 325–331 (2020)

5. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 839–847 (2018). <https://doi.org/10.1109/WACV.2018.00097>
6. Cireřan, D.C., Giusti, A., Gambardella, L.M., Schmidhuber, J.: Mitosis detection in breast cancer histology images with deep neural networks. In: International conference on medical image computing and computer-assisted intervention. pp. 411–418. Springer (2013)
7. Cree, I.A., Tan, P.H., Travis, W.D., Wesseling, P., Yagi, Y., White, V.A., Lokuhetty, D., Scolyer, R.A.: Counting mitoses: Si (ze) matters! *Modern Pathology* **34**(9), 1651–1657 (2021)
8. Galdran, A., Carneiro, G., González Ballester, M.A.: Balanced-mixup for highly imbalanced medical image classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 323–333. Springer (2021)
9. Gu, H., Liang, Y., Xu, Y., Williams, C.K., Magaki, S., Khanlou, N., Vinters, H., Chen, Z., Ni, S., Yang, C., Zhang, X.R., Li, Y., Haeri, M., Chen, X.A.: Improving workflow integration with xpath: Design and evaluation of a human-ai diagnosis system in pathology. (2021)
10. Jahanifar, M., Shepard, A., Zamanitajeddin, N., Bashir, R., Bilal, M., Khurram, S.A., Minhas, F., Rajpoot, N.: Stain-robust mitotic figure detection for the mitosis domain generalization challenge. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 48–52. Springer (2021)
11. Li, C., Wang, X., Liu, W., Latecki, L.J.: Deepmitosis: Mitosis detection via deep detection, verification and segmentation networks. *Medical image analysis* **45**, 121–133 (2018)
12. Mahmood, T., Arsalan, M., Owais, M., Lee, M.B., Park, K.R.: Artificial intelligence-based mitosis detection in breast cancer histopathology images using faster r-cnn and deep cnns. *Journal of clinical medicine* **9**(3), 749 (2020)
13. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
14. Tellez, D., Balkenhol, M., Otte-Höller, I., van de Loo, R., Vogels, R., Bult, P., Wauters, C., Vreuls, W., Mol, S., Karssemeijer, N., Litjens, G., van der Laak, J., Ciompi, F.: Whole-slide mitosis detection in h&e breast histology using phh3 as a reference to train distilled stain-invariant convolutional networks. *IEEE Transactions on Medical Imaging* **37**(9), 2126–2136 (2018). <https://doi.org/10.1109/TMI.2018.2820199>
15. Xue, C., Dou, Q., Shi, X., Chen, H., Heng, P.A.: Robust learning at noisy labeled medical images: Applied to skin lesion classification. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 1280–1283. IEEE (2019)
16. Yang, S., Luo, F., Zhang, J., Wang, X.: Sk-unet model with fourier domain for mitosis detection. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 86–90. Springer (2021)
17. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2921–2929 (2016)
18. Zhu, H., Salcudean, S., Rohling, R.: Gaze-guided class activation mapping: Leveraging human attention for network attention in chest x-rays classification. arXiv preprint arXiv:2202.07107 (2022)