

Detecting Mitoses with a Convolutional Neural Network for MIDOG 2022 Challenge

Hongyan Gu¹, Mohammad Haeri², Shuo Ni¹, Christopher Kazu Williams³, Neda Zarrin-Khameh⁴, Shino Magaki³, and Xiang 'Anthony' Chen¹

¹HCI Research, University of California, Los Angeles, USA

²Department of Pathology and Laboratory Medicine, University of Kansas Medical Center, USA

³Department of Pathology and Laboratory Medicine, UCLA David Geffen School of Medicine, USA

⁴Pathology and Immunology, Baylor College of Medicine, USA

This work presents a mitosis detection method with only one vanilla Convolutional Neural Network (CNN). Our approach 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 generalizability, we train the CNN with a series of data augmentation techniques, a loss that copes with noise-labeled images, and an active learning strategy. Our approach achieved an F1 score of 0.7323 with an EfficientNet-b3 model in the preliminary test phase of the MIDOG 2022 challenge.

Mitosis detection | domain shift | convolutional neural network | class activation map

Correspondence: xac@ucla.edu

Introduction

Mitotic activity is a crucial pathological indicator related to cancer malignancy and patients' prognosis (1). Because of its importance, a considerable amount of literature has proposed datasets (2–5) and deep learning models (6–8) for mitosis detection. State-of-the-art methods utilize a two-stage approach — a localization model (*e.g.*, RetinaNet) is employed for extracting interest locations, followed by a classification model to justify whether these locations have mitoses (4, 7, 8). Such a two-stage setup was reported to improve the performance of mitosis detection compared to that with the localization model only (4).

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 (9). Instead, our approach extracts the location of mitoses with the class activation map (CAM) (10), which allows CNNs to accept a larger input size for more efficient training. Also, our approach 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 (11). The challenge training set consists of 403 Hematoxylin & Eosin (H&E) stained 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 generalizability:

1. An augmentation pipeline with balance-mixup (12) and stain augmentation (9);
2. An Online Uncertainty Sample Mining (OUSM) (13) + COnsistent RANk Logits (CORAL) loss (14) to cope with noisy labels;
3. An active learning strategy that adds false-positive, false-negative, and hard-negative patches after each round of training.

Methods

A. Extracting Patches for Initial Training. We randomly used ~90% of the image instances in the MIDOG 2022 Challenge to generate the training set and ~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 center of each annotation (provided by the challenge) and placed them into the train/validation set.

B. Model Training. We trained an EfficientNet-b3 (15) model (input size: $240 \times 240 \times 3$) from pre-trained ImageNet weights. Here, we added model generalization 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 (9) and balance-mixup (12) – 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 and a Cosine Annealing learning rate scheduler with warm restart (max LR= 6×10^{-4}). Since we treated all unlabeled images as negative, we used an OUSM(13) + CORAL loss (14) to deal with noisy labels. Each round of training had 100 epochs, and we selected the

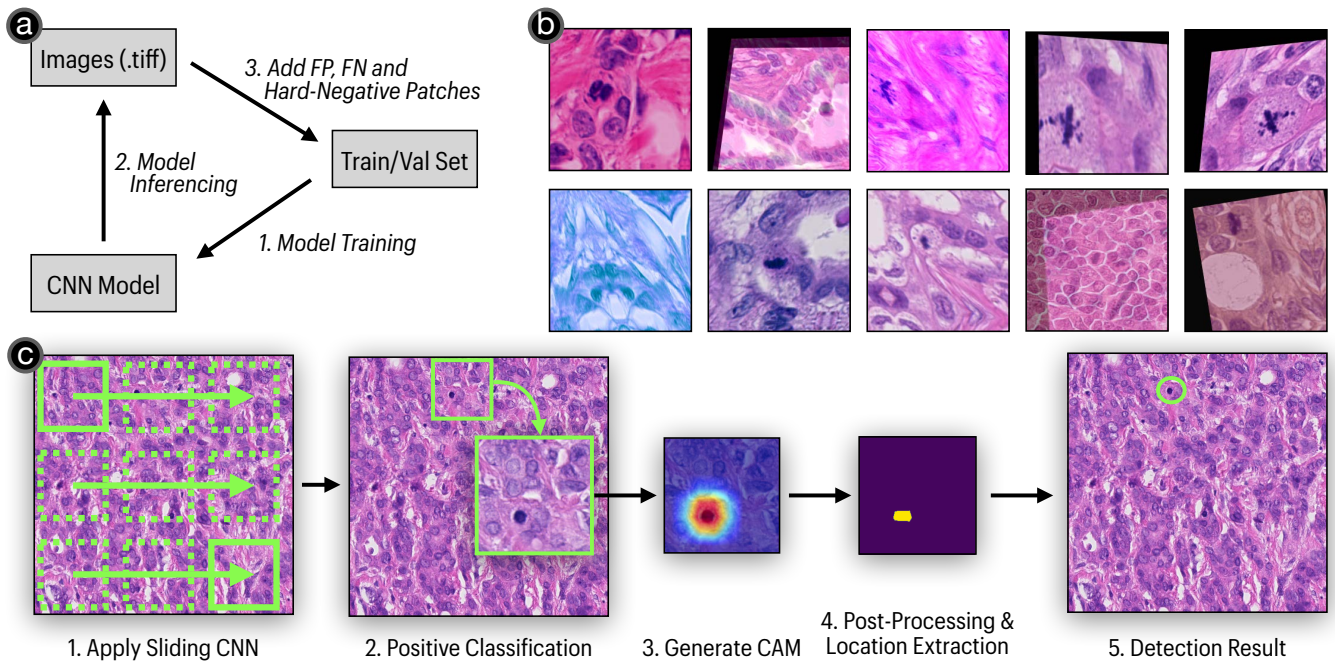


Fig. 1. (a) Illustration of the active learning training strategy used in this work; (b) Examples of augmented patches according to our augmentation pipeline; (c) Overall data processing pipeline of our approach: detecting mitosis using a convolution neural network and the class activation map.

model with the highest F1 score on the validation set for inferencing.

C. 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.

D. 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 B). Then, the best model is selected and applied to the images (Section C). 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 did not increase. Eventually, there are 103,816 patches in the final training set, and 23,638 in the validation set.

E. Extracting Mitosis Locations with CAMs. We used the best model from the final round in Section D 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 eliminate the overlapping windows. For each positive window, we calculated the CAM with GradCAM++ (16), and extracted the hotspot’s centroid as the mitosis location (Figure 1(c)).

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. More specifically, our model has F1 scores of 0.7467, 0.7593, 0.6963, and 0.7407 for detecting mitoses in four types of tumors, respectively.

Discussion

Although CAMs are primarily used for explaining CNN classifications, we demonstrate their power to detect mitoses precisely in H&E images. It is noteworthy that CAMs might fail to highlight all mitoses when there are multiple in an image. To this extent, we suggest that future work consider aligning mitosis locations and CAMs to improve detection performance and explanation quality.

Bibliography

1. Ian A Cree, Puay Hoon Tan, William D Travis, Pieter Wesseling, Yukako Yagi, Valerie A White, Dilani Lokuhetty, and Richard A Scolyer. Counting mitoses: Si (ze) matters! *Modern Pathology*, 34(9):1651–1657, 2021.
2. Roux Ludovic, Racoceanu Daniel, Loménie Nicolas, Kulikova Maria, Irshad Humayun, Klossa Jacques, Capron Frédérique, Genestie Catherine, et al. Mitosis detection in breast cancer histological images an icpr 2012 contest. *Journal of pathology informatics*, 4(1):8, 2013.
3. Mitko Veta, Yujing J Heng, Nikolas Stathonikos, Babak Ehteshami Bejnordi, Francisco Beca, Thomas Wollmann, Karl Rohr, Manan A Shah, Dayong Wang, Mikael Rousson, et al. Predicting breast tumor proliferation from whole-slide images: the tupac16 challenge. *Medical image analysis*, 54:111–121, 2019.
4. Marc Aubreville, Christof A Bertram, Taryn A Donovan, Christian Marzahl, Andreas Maier, and Robert Klopffleisch. A completely annotated whole slide image dataset of canine breast cancer to aid human breast cancer research. *Scientific data*, 7(1):1–10, 2020.
5. Christof A Bertram, Marc Aubreville, Christian Marzahl, Andreas Maier, and Robert Klopffleisch. 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.
6. Dan C Cireşan, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber. Mitosis detection in breast cancer histology images with deep neural networks. In *International conference on medical image computing and computer-assisted intervention*, pages 411–418. Springer, 2013.

7. Chao Li, Xinggong Wang, Wenyu Liu, and Longin Jan Latecki. Deepmitosis: Mitosis detection via deep detection, verification and segmentation networks. *Medical image analysis*, 45:121–133, 2018.
8. Tahir Mahmood, Muhammad Arsalan, Muhammad Owais, Min Beom Lee, and Kang Ryoung Park. 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.
9. David Tellez, 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 Ciompi. Whole-slide mitosis detection in h&e breast histology using ph3 as a reference to train distilled stain-invariant convolutional networks. *IEEE Transactions on Medical Imaging*, 37(9):2126–2136, 2018. doi: 10.1109/TMI.2018.2820199.
10. Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.
11. Marc Aubreville, Christof Bertram, Katharina Breining, Samir Jabari, Nikolas Stathonikos, and Mitko Veta. Mitosis domain generalization challenge 2022. In *25th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2022)*, 2022. doi: 10.5281/zenodo.6362337.
12. Adrian Galdran, Gustavo Carneiro, and Miguel A González Ballester. Balanced-mixup for highly imbalanced medical image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 323–333. Springer, 2021.
13. Cheng Xue, Qi Dou, Xueying Shi, Hao Chen, and Pheng-Ann Heng. Robust learning at noisy labeled medical images: Applied to skin lesion classification. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1280–1283. IEEE, 2019.
14. Wenzhi Cao, Vahid Mirjalili, and Sebastian Raschka. Rank consistent ordinal regression for neural networks with application to age estimation. *Pattern Recognition Letters*, 140: 325–331, 2020.
15. Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
16. Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 839–847, 2018. doi: 10.1109/WACV.2018.00097.