Cancer Diagnosis – PCam

Simon Zou, Yehchan Yoo

Introduction

Diagnostic Challenge: Accurate identification of metastatic tissue in lymph node histopathological scans is essential for cancer diagnosis.

Pathologist Shortage: A severe shortage of pathologists is straining healthcare systems worldwide. In the United States alone, the number of active pathologists dropped by nearly 20% between 2007 and 2017, while their workload surged by over 40%. This trend remains far worse in low- and middle-income countries with far less resources.

Project Goal: Our project aims to harness deep learning to automate tumor identification in histological image patches using the Patch Camelyon (PCam) dataset. The use of artificial intelligence (AI) for tasks like tumor detection can help a lot in reducing the workload of the existing oncologists and allowing them to focus on more complex tasks.

Baseline G-CNN for PCam

Multi-Directional Filtering: Applies rotated and reflected versions of a filter simultaneously to detect patterns (tumor cells) in any orientation from the start.

Tracking Orientation: Uses "grouped" feature maps throughout the network to record both what feature was found and which way it was facing.

Combining for Consistency: Mathematically combines these orientation-aware maps at the final layer, guaranteeing a consistent prediction even if the input image is spun or mirrored.

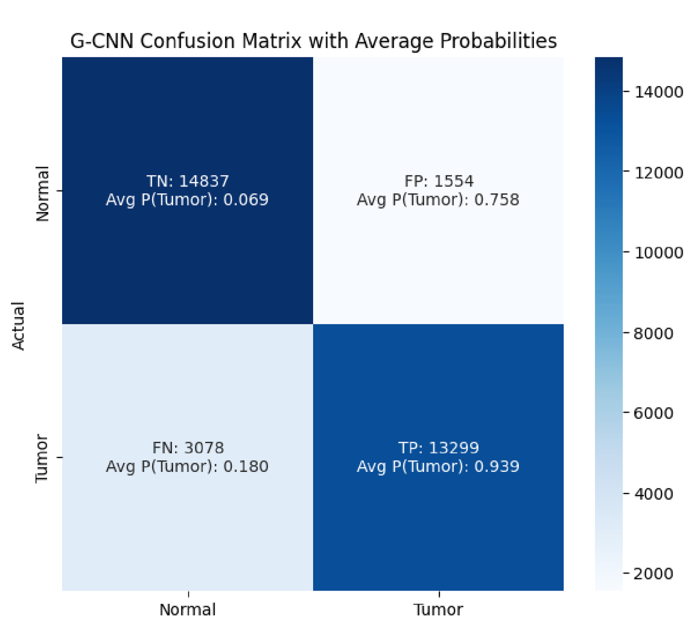

Results:
Accuracy: 0.8586
AUC-ROC: 0.9301

Tumor precision: 0.90
Normal precision: 0.83
Tumor recall: 0.81
Normal recall: 0.91
Tumor f1-score: 0.85
Normal f1-score: 0.86

Method 1: Custom CNN

Feature Extraction: Three sequential convolutional blocks (32, 64, and 128 filters) progressively learn simple to complex cellular patterns.

Stabilization & Downsampling: Each block uses Batch Normalization to stabilize training and Max-Pooling to shrink the image footprint (from 96x96 down to 12x12).

Binary Classification: A fully connected head uses a 512-neuron layer with Dropout (to prevent overfitting) and outputs a single probability for tumor presence.

4 trials were conducted with different parameters:

1. Batch size: 64, Epochs: 10, Learning rate: 0.001
   - Accuracy: 0.8043, AUC-ROC: 0.9055
2. Batch size: 32, Epochs: 20, Learning rate: 0.001
   - Accuracy: 0.8074, AUC-ROC: 0.9132
3. Batch size: 64, Epochs: 10, Learning rate: `ReduceLROnPleateau`
   - Accuracy: 0.8391, AUC-ROC: 0.9148
4. Batch size: 64, Epochs: 10, Learning rate: `ReduceLROnPleateau` `With data augmentation`

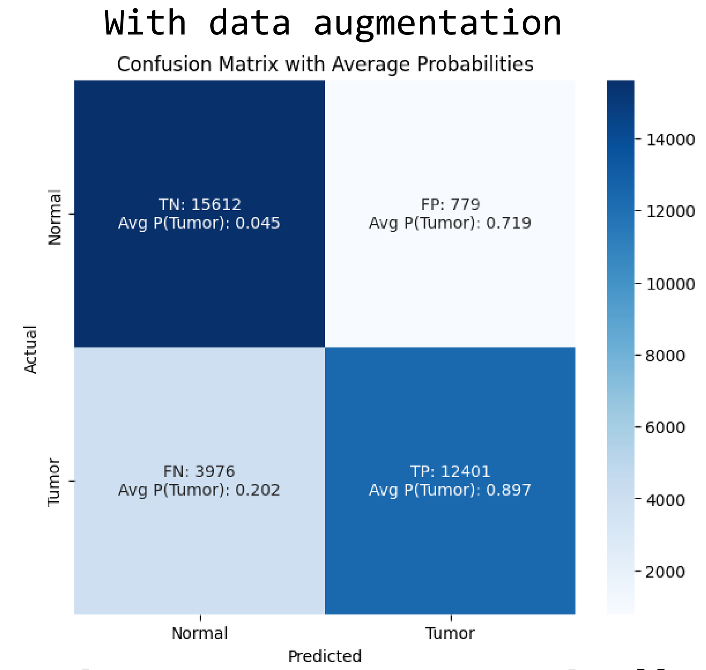

Trial 4 Results:
Accuracy: 0.8549
AUC-ROC: 0.9464

Tumor precision: 0.94
Normal precision: 0.80
Tumor recall: 0.76
Normal recall: 0.95
Tumor f1-score: 0.84
Normal f1-score: 0.87

Shapley Values Plotted

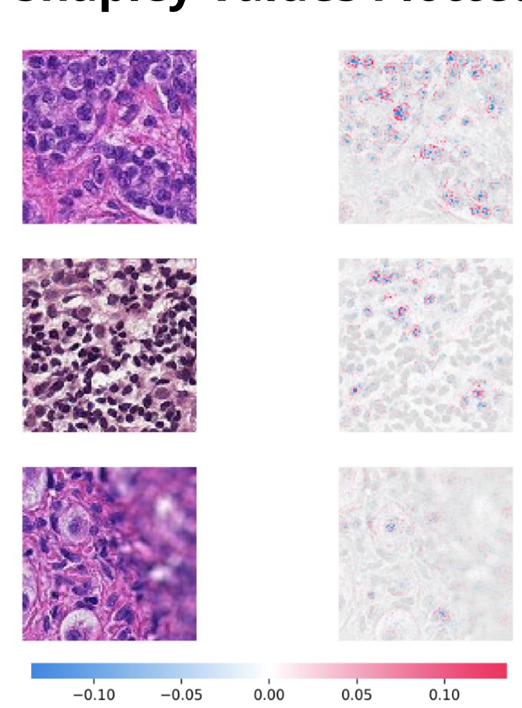

Heatmap using Grad-CAM

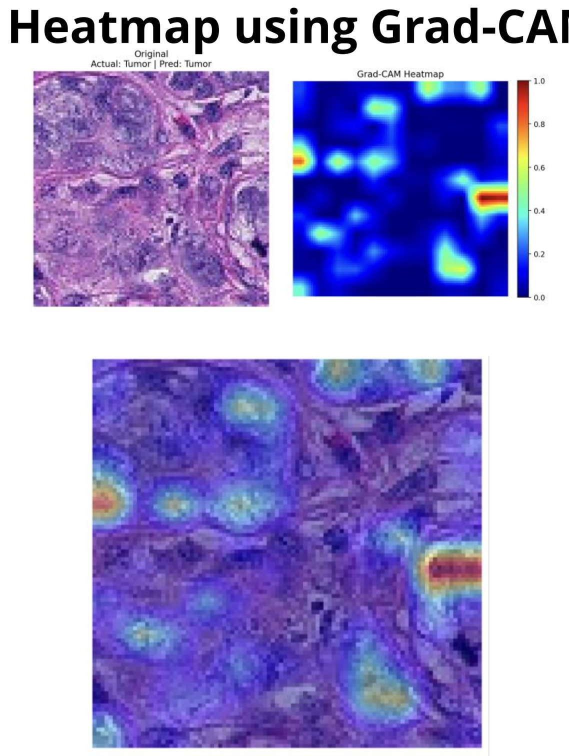

Method 2: ResNet18

Pre-Trained CNN Model: 18-layer convolutional neural network model developed and released by Microsoft Research in 2015

Residual Learning: Approximates a residual function for the underlying mapping function with identity mappings

Fine-Tuning: Fine-tuned with hyperparameter tuning on the train PCam dataset with Nvidia P100 GPU on Kaggle using PyTorch and OpTuna

Version 1 of ResNet18 Model

Hyperparameters: Learning rate of ~0.00017, batch size of 128, Adam optimizer

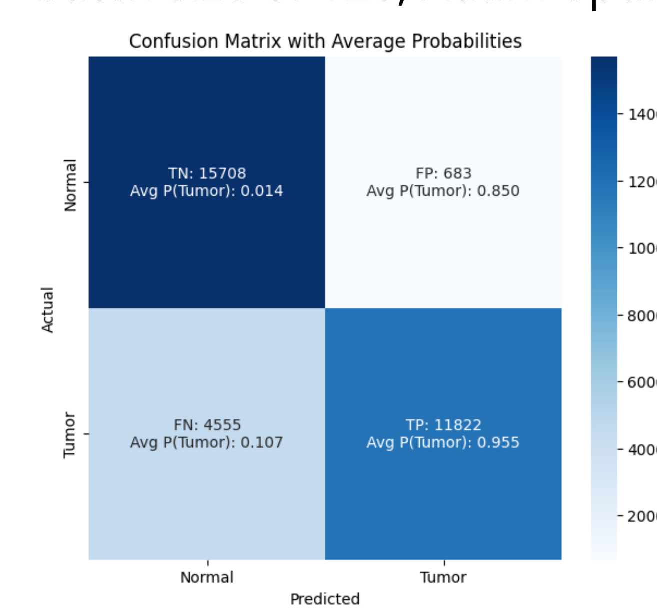

Results:
Accuracy: 0.8401
AUC-ROC: 0.9418

Tumor precision: 0.95
Normal precision: 0.78
Tumor recall: 0.72
Normal recall: 0.96
Tumor f1-score: 0.82
Normal f1-score: 0.86

Shapley Values Plotted

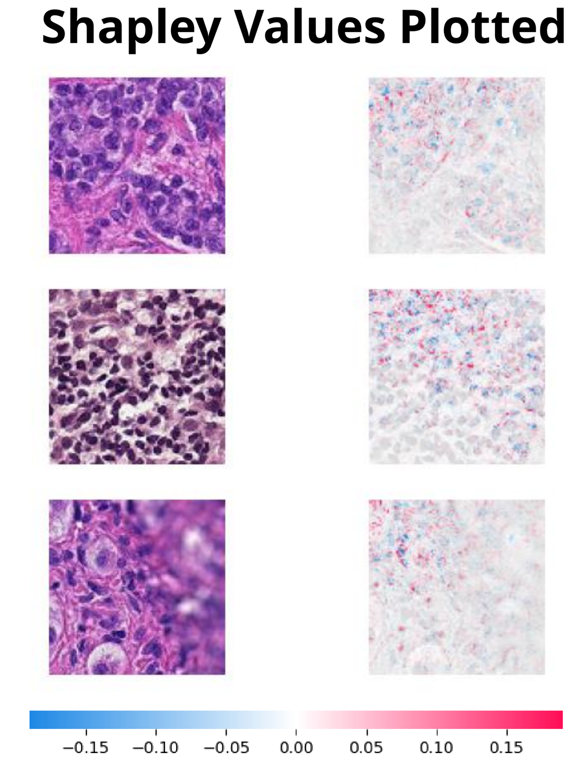

Heatmap using Grad-CAM

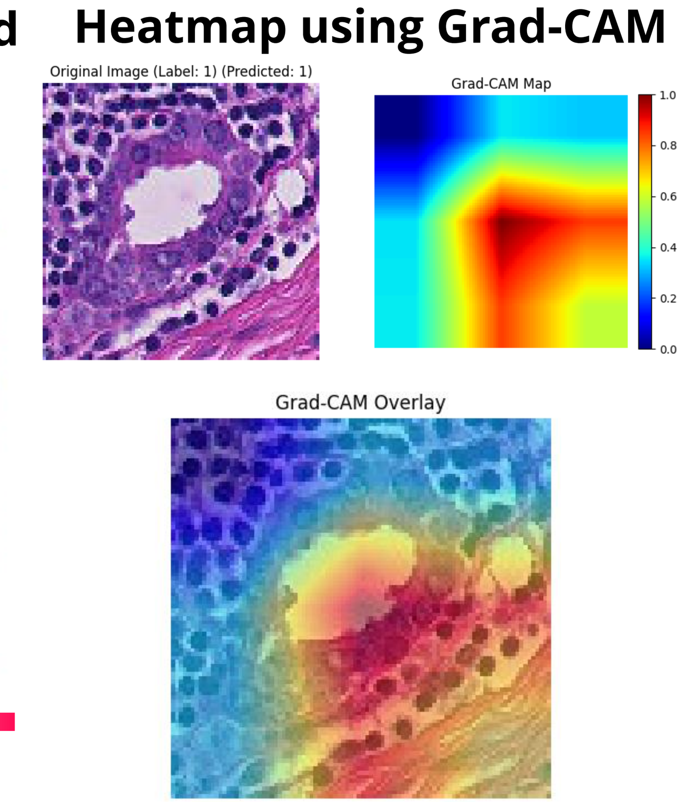

Version 2 of ResNet18 Model

Data Augmentation: Randomly applied the following transformations to the training set:

- Horizontal flip
- Vertical flip

Data Augmentation (cont.):

- Rotation by up to 90 degrees either direction
- Brightness changed up or down by up to 20%
- Contrast changed up or down by up to 20%
- Saturation changed up or down by up to 20%
- Hue changed up or down by up to 5%

Hyperparameters: Learning rate of ~0.00286, batch size of 32, SGD optimizer

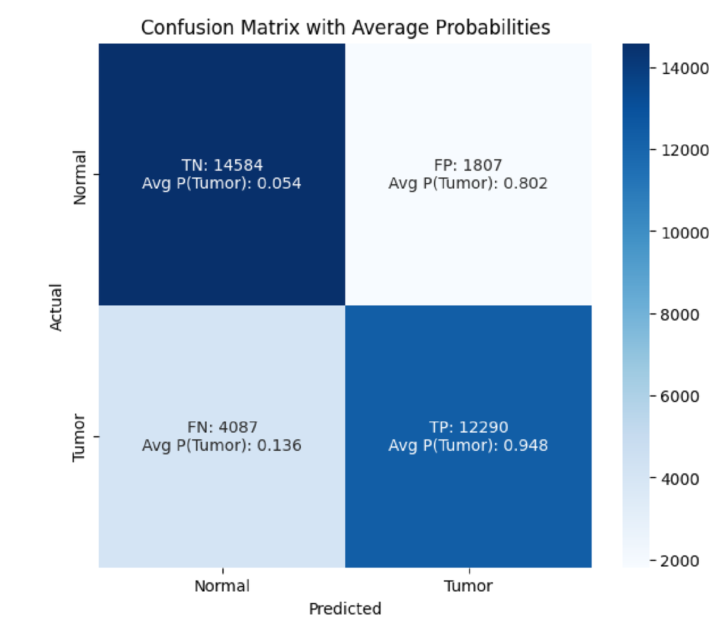

Results:
Accuracy: 0.8665
AUC-ROC: 0.9441

Tumor precision: 0.93
Normal precision: 0.82
Tumor recall: 0.79
Normal recall: 0.94
Tumor f1-score: 0.86
Normal f1-score: 0.88

Shapley Values Plotted

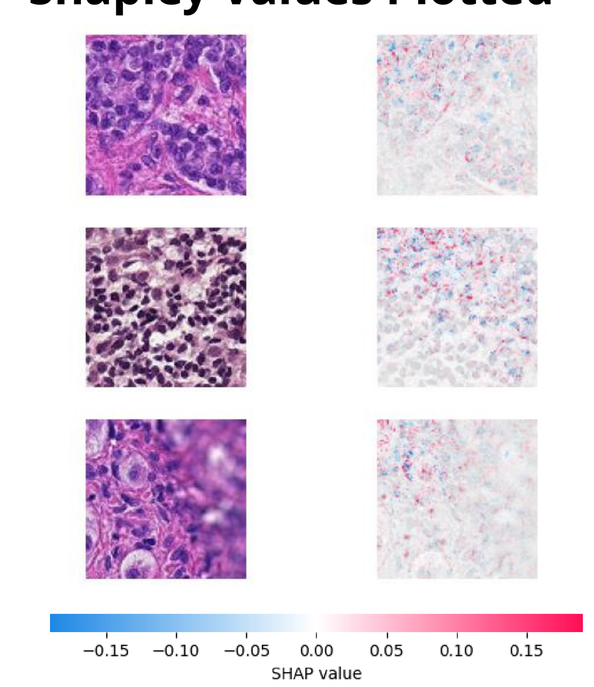

Heatmap using Grad-CAM

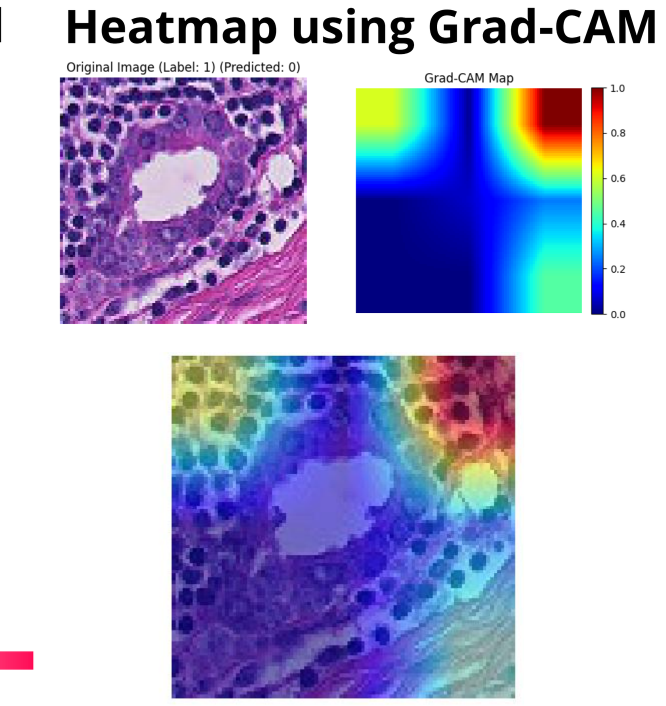

Carbone FG. The Shrinking Workforce of Pathologists: Implications for Healthcare and Possible Solutions. Pathologica. 2025;117(4):449-451. doi:https://doi.org/10.32074/1591-951x-n1156

Oluwaseun Adebayo Bamodu, Chung CC. Cancer Care Disparities: Overcoming Barriers to Cancer Control in Low- and Middle-Income Countries. JCO Global Oncology. 2024;10(10). doi:https://doi.org/10.1200/go.23.00439

PatchCamelyon - Grand Challenge. Grand Challenge. Published 2017. Accessed January 23, 2026. https://patchcamelyon.grand-challenge.org/

Veeling B, Linmans J, Winkens J, Cohen T, Welling M. Rotation Equivariant CNNs for Digital Pathology. Published online June 8, 2018. Accessed February 19, 2026. https://arxiv.org/abs/1806.03962

Veeling B. PatchCamelyon (PCam). GitHub. Published January 31, 2024. Accessed February 20, 2026. https://github.com/basveeling/pcam

GeeksforGeeks. Introduction to convolution neural network. GeeksforGeeks. July 11, 2025. Accessed February 19, 2026. https://www.geeksforgeeks.org/machine-learning/introduction-convolution-neural-network/.

Training parameters - amazon machine learning. Accessed February 20, 2026. https://docs.aws.amazon.com/machine-learning/latest/dg/training-parameters1.html.

He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. Published online December 10, 2015. Accessed February 19, 2026. https://arxiv.org/abs/1512.03385

resnet18. PyTorch. Accessed February 19, 2026. https://docs.pytorch.org/vision/main/models/generated/torchvision.models.resnet18.html

Najib T. Hyperparameter Tuning Using Optuna. Medium. Published August 11, 2023. Accessed February 19, 2026. https://medium.com/@taeefnajib/hyperparameter-tuning-using-optuna-c46d7b29a3e

Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-cam: Visual explanations from deep networks via gradient-based localization. arXiv.org. December 3, 2019. Accessed February 19, 2026. https://arxiv.org/abs/1610.02391.