SIIM CMIMI24
Oct 21-22 | Boston, MA
Conference On Machine Intelligence In Medical Imaging

Two-Step Fully Automated Classification of Choroidal Metastases on MRI: Orbit Localization via Bounding Boxes Followed by Binary Classification via Evolutionary Strategies

Joseph N. Stember, MD, PhD, Memorial Sloan Kettering Cancer Center; Bala McRae-Posani, MSc; Andrei Holodny, MD
Hrithwik Shalu, MS; Jeffrey S. Shi

Introduction/Background

The choroid of the eye is a rare site for metastatic spread of a tumor, and choroidal metastases (CMs) may be visualized on magnetic resonance imaging (MRI). However, as small lesions on the periphery of the image, they are often missed on brain MRI.

Methods/Intervention

Here, we describe sequential cropping and classification on brain MRI images to detect CMs using artificial intelligence (AI). We first trained an orbit localization model with a YOLOv5 architecture using 386 normal T2-weighted brain MRI images. The model predicted and cropped the positions of the orbits on MRI brain scans from 33 patients without and 33 patients with CMs. After zooming in around the orbits, the cropped images served as inputs to a binary classifier convolutional neural network (CNN) to classify images as normal or CM-containing. We used 36 images for training and the other 30 for testing. Given the small training set, we trained the network weights via the data-efficient deep neuroevolution (DNE) strategy.

Results/Outcome

Our orbit localization model achieved mean average precision at intersection over union of 0.5 of 0.590. For a confidence of 0.3, the model achieved recall of 1.00 and precision of 0.50, as the model accurately identified all orbits but was unable to distinguish “left” and “right” (Figure 1). Laterality was assigned afterwards using relative position. The model generalizes to scans with CMs; on our dataset of 33 slices demonstrating CMs, the model accurately determined the bounding boxes without errors (Figure 2). The predicted bounding boxes were used to crop the images for training our CNN classification model. After training via DNE for over 80,000 episodes, the model converged on a training set accuracy of 100% and testing set accuracy of 100% (Figure 3).

Conclusion

We trained a YOLOv5 model to accurately localize and crop the orbits on brain MRI. The cropped images were subsequently used to train a CNN with excellent performance in detecting CMs.

Statement of Impact

Our method provides an end-to-end model to accurately detect small, peripheral, easy-to-miss lesions to potentially improve sensitivity for detection of CMs. It could thereby help reduce “corner of the image” false negatives.

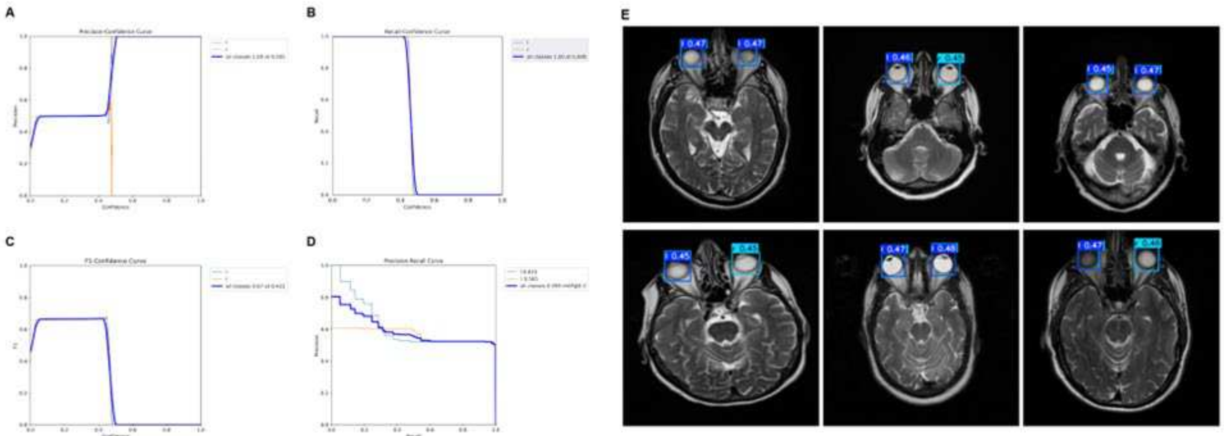

Figure 1: Performance of orbit localization model. A: Precision-confidence curve. B: Recall-confidence curve. C: F1-confidence curve. D: Precision-recall curve. Discontinuities at confidence=0.5 reflect the inability of the model to distinguish laterality. E: Representative images of model output predictions on normal brain MRIs.

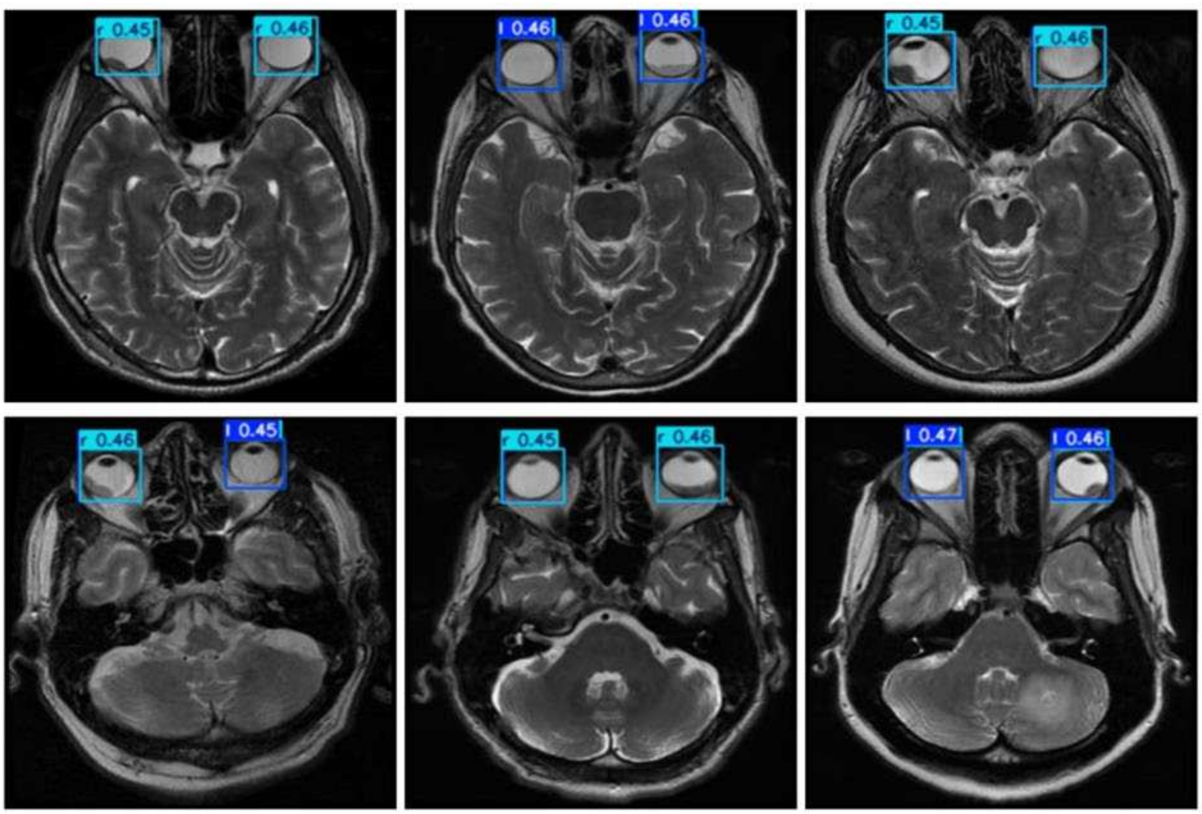

Figure 2: Representative images demonstrating the performance of our orbit localization model on brain MRIs with choroidal metastases.

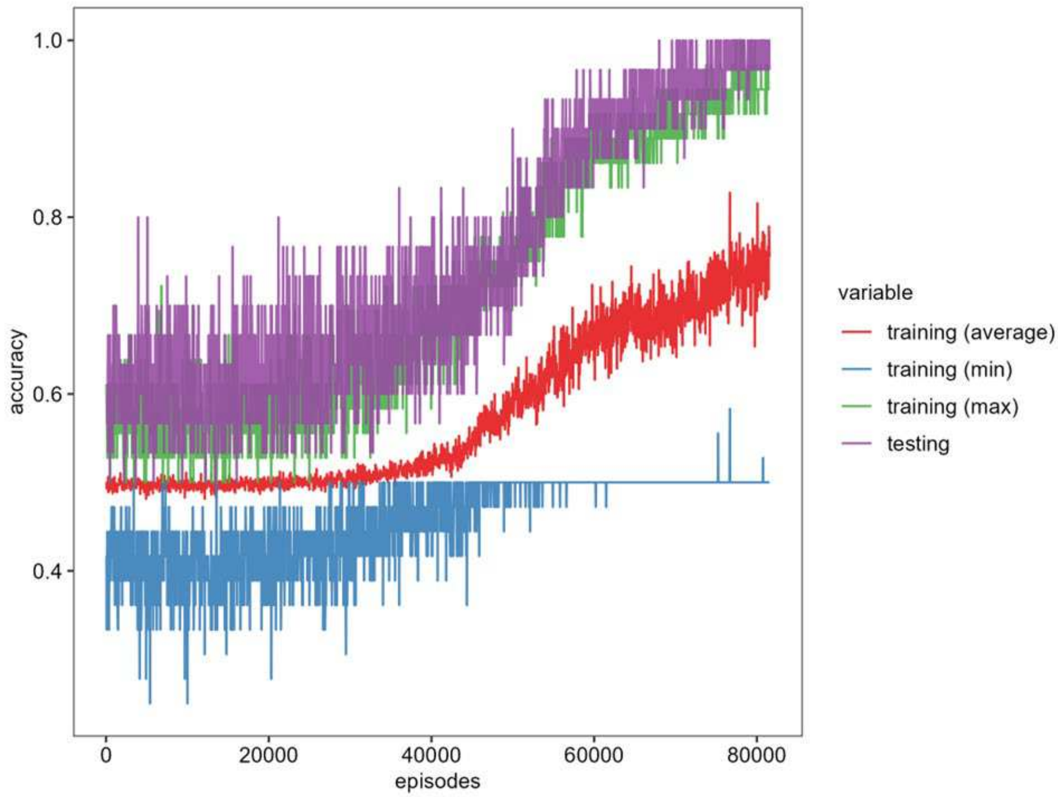

Figure 3: CNN accuracy versus episodes during training. The average, worst, and best performing "children" of the DNE strategy during training are represented as training (average), training (min), and training (max), respectively.

Keywords

object detection**;** classification**;** tumor**;** cancer**;** metastasis