Detection of Melanoma via Deep Learning

Andrews, Sophie (School: Chatham High School)

Convolutional Neural Networks (CNNs) are machine learning systems that can classify images. Applying CNNs to medical images for recognizing different types of skin lesions can ultimately create a more convenient and affordable way to diagnose skin cancer. I use Caffe, a deep learning platform, to train a standard image classification CNN on over 2000 melanoma (cancerous) and nevus (benign) images from the International Skin Imaging Collaboration (ISIC). I feed the images along with their rotational variants through the CNN and then Caffe adjusts weights within the CNN to optimize classification accuracy. I call the resulting network SkinNet. SkinNet has comparable classification performance to the best results in the literature. In testing a separate set of over 1000 lesion images, SkinNet obtains an “Area Under the Curve” (AUC) of 0.869, where AUC measures the tradeoff between the true positive and true negative rates. I also observe that SkinNet is more likely to err on melanoma images that do not have the textbook ABC-characteristics of melanoma. I am currently building a website that allows a user to upload a smartphone image and obtain the most likely melanoma/nevus classification as well as the most similar image from the ISIC database. I conclude that CNNs for general purpose image recognition can be trained for the special task of classifying melanoma and nevus. Moreover, SkinNet responds to similar lesion characteristics that humans use for melanoma detection.

Awards Won: