Utilizing Machine Learning Techniques to Identify Cancerous Skin Lesions

Kang, Christopher

Skin cancer is the most common form of cancer. When diagnosed early, there is a 98% survival rate. However, once the cancer metastasizes, the survival rate decreases dramatically to 18%. The objectives of this study were to (a) use machine learning techniques to better identify the behavior of moles and accelerate the diagnosis of these skin lesions and (b) to minimize the computational cost of training these techniques. The project included assembling the datasets, building a simple feedforward neural network, and applying different convolutional neural networks. Additionally, I assembled a workstation to accelerate the neural networks. My project compared simple feedforward neural networks (FNN) to convolutional neural networks (CNN) in a melanoma image classification task. I programmed my neural networks in Python using Tensorflow and NVIDIA DIGITS, respectively. The networks were trained using images subset of the International Skin Imaging Collaboration dataset (ISIC). The ISIC database has 2000 images that have been collected by dermatologists in an effort to standardize dermatological images. Three different types of datasets used in machine learning: training, validation, and testing. The results show the Alexnet CNN outperformed the simple FNN because of its inherent complexity and ability to process images more effectively, at the cost of computational resources. The simple FNN, which can be trained more quickly, received about 70% accuracy. The CNN, which required a GPU to accelerate, achieved an 85% accuracy which was better than available programs and close to the 88% accuracy by diagnosis through a dermatologist.

Awards Won:

Fourth Award of $500