My Skin: A Deep Convolutional Neural Model for Skin Cancer Identification with Coarse-to-Fine Contextual Memory (CFCM)

Jadhav, Ameya (School: Denmark High School)

Skin cancer, the most common human malignancy, is diagnosed visually, beginning with an initial clinical screening and followed with a dermoscopic analysis, a biopsy and histopathological examination. An intelligent, accurate automation of skin lesion identification and classification is a challenging task, mainly because of the fine-grained variability in the appearance of skin lesions. Deep convolutional neural networks (CNNs) show high potential for general and highly variable tasks across many fine-grained object categories. This project utilizes deep CNNs, trained end-to-end from images and disease labels as inputs. The CNN is trained on 10,682 clinical images. Its performance has been tested against an additional 3,561 clinical test images and 3,562 clinical validation images for two critical binary classification use cases: melanoma, or not melanoma. This project explores various machine learning models and training hyperparameters, finetuned to achieve the highest binary validation accuracy. Further, Coarse-to-Fine Contextual Memory (CFCM) has been used for deep encoder-decoder architecture clinical image segmentation to improve overall performance. Various other high utility machine learning algorithms and methods have been used, including data pipelines, unique image augmentation, and transfer learning. The project achieves performance on par with board-certified dermatologists (92.7%), demonstrating an artificial intelligence capable of classifying skin cancer with a level of competence comparable to dermatologists. This project also covers a developed iOS app, named My Skin, to distribute the use of the artificial intelligence (among other features), and can therefore potentially provide low-cost universal access to vital diagnostic care.