Deep Learning-Based Oral Lesion Classification Leveraging Real-Time 3D Gradient-Weighted Class Activation Mapping (Grad-CAM) Imaging

Reddy, Sumedh (School: Central High School)
Oghlakian, Yohann (School: Central High School)

Each year, about 58,500 Americans are diagnosed with oral or oropharyngeal cancer, and over 12,000 of them will not survive it. When oral cancer is caught at a localized stage, the 5-year survival rate is around 83.7%, but if it has already spread by the time it is found, it drops to 38.5%. Early detection is one of the most important factors in surviving this disease, yet current screening methods require in-person visits to a specialist, which is not always accessible. This project aims to help address that problem by developing a system that can detect oral cancer and other oral abnormalities from images. To build this system, we collected over 10,000 clinical images through partnerships with Washington University and Leland Stanford Junior University, along with public datasets. Using this data, we trained a custom convolutional neural network (CNN) to classify oral images into different categories based on the type of abnormality present. Building our own CNN rather than using a pre-existing model allowed us to design the architecture specifically around oral pathology. Our model achieved a high classification accuracy of 98%, outperforming models such as DenseNet (~90%) and EfficientNet (~97%). The biggest challenge with using machine learning in medicine is that models are considered "black boxes" because they answer without explanation. To address this, we integrated Grad-CAM (Gradient-weighted Class Activation Mapping). It generates a heatmap that shows which areas of the image the model focused on, making the system more transparent for medical professionals to trust in practice.

Awards Won: