Skin Cancer Detection Executive Summary

Objective: To develop AI algorithms that accurately identify skin cancer using 3D total body photographs and patient metadata.

Team Members:

Madelyn Esther Cruz, **Maksim Kosmakov**

Overview

This project aims to enhance early skin cancer detection through a machine learning approach that integrates both image and tabular data. Early identification of melanoma, basal cell carcinoma, and squamous cell carcinoma is critical for improving patient outcomes.

Key Achievements

- Developed and tested 51 models on Kaggle, with a focus on optimizing the Partial Area Under the ROC Curve (pAUC) above an 80% True Positive Rate (TPR).
- The best model, trained on an oversampled dataset (100k benign images and 10k melanoma images), achieved a pAUC of 0.140, reflecting strong diagnostic performance.

Methodology

- **Data:** SLICE-3D dataset with patient demographics, lesion characteristics, and dermoscopy images.

- **Models:** ResNet50 and EfficientNetV2, combined into an ensemble model fine-tuned with CrossEntropy loss, L2 regularization, and the Adam optimizer.

Future Directions

- Improving image preprocessing to remove artifacts.
- Exploring additional ensemble techniques.
- Managing high memory requirements for model training.

Conclusion:

This project has demonstrated the potential of AI in aiding early skin cancer detection, achieving promising results that warrant further refinement and exploration.