University Ph.D. Dissertation Defense
Department of Electrical Engineering

Philip M. Adamson

Medical Image Acquisition and Representation Learning for Precision Oncology

Advisor: Daniel Spielman

Date: Friday, December 12, 2025
Time: 1:00 pm
Location: AllenX 101X
Zoom Link: https://stanford.zoom.us/j/94764811232?pwd=bfPF9oAJbliyEvRCN73DENqINYJnUU.1
Password: 844774

Advances in medical imaging depend on the ability to acquire informative biological features, assess the diagnostic fidelity of reconstructed images, and learn image feature representations that support downstream clinical tasks. This dissertation presents methods spanning these three components of the medical imaging pipeline, from acquisition to analysis, in the context of precision oncology.

Metabolic reprogramming is a hallmark of cancer and a target of emerging therapeutics, motivating the need for clinically translatable imaging methods that directly quantify metabolic activity in vivo. Deuterium Metabolic Imaging (DMI) maps downstream metabolic products of deuterated glucose, offering a potential biomarker for monitoring therapeutic response in oncology**.** In this work, I describe how to enable DMI at the widely available 3T MRI field strength and demonstrate whole brain 3T DMI in patients with central nervous system lesions.

Beyond acquiring novel metabolic imaging features, the quality of conventional MR images remains essential for extracting reliable biomarkers in oncology. However, accelerated MR reconstructions are often evaluated with metrics that poorly reflect diagnostic utility. To address this, I develop deep feature distances (DFDs), fully automated data-driven image quality metrics, for assessing MR image reconstruction quality. In a blinded radiologist reader study, DFDs correlate more strongly with diagnostic utility than SSIM and PSNR, positioning them as more clinically relevant tools for evaluating the fidelity of diagnostically important image features.

Building on these improvements in imaging acquisition and evaluation, I develop MerlinOnc, a vision-language CT foundation model for thoracic oncology. Trained on over 15,000 chest CT scans and linked clinical data from 4,500 thoracic oncology patients, MerlinOnc's learned image features outperform other state-of-the-art CT foundation models on tasks such as cancer staging, progression assessment, and survival prediction. These contributions illustrate how advances in medical image acquisition, evaluation, and representation learning can yield more clinically impactful imaging features for precision oncology.