University of Pittsburgh | School of Computing and Information

Proposal Defense

Doctor of Philosophy in Intelligent Systems Program

"**Evidence-Preserving AI for Longitudinal and Multimodal Breast Cancer Assessment**"
by **Zhengbo Zhou**

Date: November 6, 2026
Time: 2 to 4 p.m.
Place: 6106 Sennott Square, 210 S. Bouquet Street, Pittsburgh, PA 15213

Committee:

- Shandong Wu, PhD (Advisor), Professor, Intelligent Systems Program, University of Pittsburgh
- Leming Zhou, PhD, Professor, Intelligent Systems Program, University of Pittsburgh
- Yufei Huang, PhD, Professor, Intelligent Systems Program, University of Pittsburgh
- Lu Tang, PhD, Professor, School of Public Health, University of Pittsburgh

Abstract:
Artificial intelligence has the potential to improve breast cancer risk prediction and diagnosis by extracting evidence information from multimodal imaging and clinical data. However, clinically meaningful evidence may be subtle, localized to small anatomical regions, or apparent only through in-depth assessment across examinations, imaging modalities, and clinical context. Current AI models lack explicit mechanisms to identify, integrate, and preserve such evidence to maintain predictive performance and clinical interpretability. This dissertation develops evidence-preserving AI methodology that places anatomical, temporal, and clinical relationships at the center of model design and evaluation for breast cancer, through four connected pillars: longitudinal representation learning for breast cancer risk prediction, longitudinal and multimodal representation learning for breast cancer diagnosis, knowledge-based breast imaging AI design, and adversarial safety to protect integrity of evidence in longitudinal AI models. The research develops models that accommodate variable histories of screening mammograms and irregular examination intervals, capture bilateral, multi-view, and temporal differences, and connect localized imaging findings with clinical knowledge and radiologists' assessments. Clinically informed supervision guides these representations toward task-relevant features, supporting a unified approach to assess longitudinal cancer risk and conduct diagnostic interpretation. The same principle extends to breast MRI synthesis, where clinically useful structure and information are retained through anatomy-aware techniques in generating synthetic images. In addition, to preserve the integrity of the evidence, the research examines how adversarial perturbations disrupt longitudinal relationships and evaluates the protection provided by robust training. Together, these studies contribute novel methods for specifying supervision, preserving anatomical and temporal structure, integrating diagnostic evidence, and enhancing cyber robustness against adversaries. This dissertation advances trustworthy breast imaging AI model design and evaluation through preserving clinically grounded evidence for breast cancer risk assessment and diagnosis.