

## **Ph.D. DISSERTATION DEFENSE**

**Candidate:** Yifeng Peng  
**Degree:** Doctor of Philosophy  
**School/Department:** Charles V. Schaefer, Jr. School of Engineering and Science /  
Department of Systems Engineering  
**Date:** October 28, 2026  
**Time/Location:** 11:00 am / Babbio Center 503  
**Title:** QUANTUM FOR AI AND AI FOR QUANTUM: TOWARD  
FUTURE HIGH-PERFORMANCE HYBRID COMPUTING

**Chairperson:** Ying Wang : Charles V. Schaefer, Jr. School of Engineering and  
Science / Department of Systems Engineering

Philip Odonkor: Charles V. Schaefer, Jr. School of Engineering  
and Science / Department of Systems Engineering

**Committee Members:** Ting Liao: Charles V. Schaefer, Jr. School of Engineering and  
Science / Department of Systems Engineering

Yuping Huang: Charles V. Schaefer, Jr. School of Engineering  
and Science / Department of Physics

## **ABSTRACT**

Artificial intelligence (AI) and quantum computing offer two directions for designing hybrid computational systems: quantum circuits can provide task-specific components within AI models, while AI methods can help determine how quantum circuits are initialized, trained, and compiled. Realizing this reciprocal relationship requires connecting the quality of a quantum component's contribution to the resources needed to obtain it. This dissertation investigates that connection through six studies organized around Quantum for AI and AI for Quantum, and develops a co-design framework for future high-performance hybrid computing.

The Quantum for AI studies examine where quantum computation can contribute within an established learning architecture. Quantum squeeze-and-excitation (QSE) uses a variational circuit for feature-channel recalibration, and Qsco uses one for open-set anomaly scoring. Both retain a classical backbone for feature extraction and assign the quantum module a specific statistical role. In the evaluated simulations, linear QSE-ResNet-18 reaches 85.00% accuracy on CIFAR-10, compared with 82.20% for classical SE-ResNet-18; depth-two Qsco improves its matched baseline on seven of eight one-example general-setting datasets. These results support targeted quantum module design, with benefits that depend on circuit topology, depth, and task setting.

The AI for Quantum studies address the training and execution costs that constrain such a computing paradigm, using variational quantum eigensolvers as the experimental testbed. Two initialization studies move from transferring classical deep-learning heuristics to learning from quantum-

objective feedback: direct adaptations of Xavier, LeCun, He, and orthogonal initialization yield limited and inconsistent benefits, whereas reinforcement learning finds more favorable starting regions in selected settings. TITAN learns from prior optimization trajectories to predict parameters that can remain fixed in a new run, reducing recurring gradient evaluations. FALCON combines learned hypergraph scoring, objective saliency, and Clifford-oriented rewriting to select a compact circuit before repeated optimization and evaluate it after hardware-aware mapping. Separate evaluations report up to 10.2 times fewer parameters and a 4.0 times improvement in estimated time to a target energy.

Together, the studies motivate an architecture in which classical AI extracts representations and guides quantum resource allocation, quantum modules perform selected task operations, and task and execution feedback inform subsequent design decisions. The framework connects application quality, trainability, and execution efficiency through explicit accounting for learning, compilation, and repeated circuit evaluation. The evidence comes from simulation and hardware-aware mapping; the complete architecture and transfer of the VQE methods to the AI modules remain future work. The dissertation thus contributes evaluated components and design principles toward a computing paradigm in which AI helps make quantum computation practical and quantum computation expands the design space of AI.