Van Tien Nguyen

Ph.D. Researcher in Quantum Machine Learning, Multilinear Learning, and Tensor Methods

tien.nguyen@utsa.edu | +1 (210) 672-5784 | San Antonio, TX

vantnprof.github.io | linkedin.com/in/vantnprof | ORCID: 0009-0000-8834-9415

Research Profile

Ph.D. researcher developing structure-aware and resource-efficient methods for quantum, hybrid quantum–classical, and deep learning systems. My work preserves tensor and low-rank structure across model design, measurement, optimization, and deployment, with the goal of reducing parameter, data-processing, and quantum-resource costs while retaining predictive performance. I evaluate methods through theory, reproducible simulation, and experiments on IBM Quantum hardware using PyTorch, Qiskit, and PennyLane.

Research Interests

My research starts from a simple frustration: most hybrid quantum–classical models throw away the structure already sitting in their data, flattening images, sequences, and quantum measurements into vectors before a single parameter is learned, and then spend a quantum circuit's scarce qubits and shots relearning what was already there. In *Quantum-Based Tensor Contraction Layers* (QTCL, SPIE 2026), I addressed this by placing a small trainable quantum circuit inside a compressed tensor bottleneck rather than the full image, tying its cost to the bottleneck width instead of the raw input. The design beat its classical counterpart on VGG-19/CIFAR-100 and, just as usefully, lost to it on AlexNet — the kind of regime-dependent result that turns "does quantum help" into the more answerable question of *when* it does. That same instinct carried into purely classical territory in *Multilinear Transformation Layers* (MTL, Asilomar 2026) and its adaptive-rank extension, *LR-MTL* (under review, IEEE TAI), where late convolutions are replaced by mode-wise transformations and effective rank is learned rather than fixed by hand, and into optimization itself in *QALR-CPD*, where I formulated adaptive CP-rank selection as an interaction-aware QUBO so that a quantum solver decides which tensor components actually matter, rather than pruning by magnitude alone.

A well-placed quantum circuit is still only as useful as what you choose to measure from it, and most of the measurement-efficiency literature treats that choice as fixed rather than as part of the model. In *Parameter-Efficient Multilinear Readout for Quantum Reservoir Computing* (ML-QRC, QML@QCE 2026), I preserved the time–qubit–observable structure of reservoir measurements instead of flattening them, cutting readout parameters by roughly 12.1× without losing predictive quality. In *Variational Quantum Bayesian Regression* (VQBR, IEEE QCE 2026), I rebuilt a classical regression problem around exactly what a shallow circuit can measure, separating the target vector into a direction recovered by the circuit and a scale solved for analytically, recovering the true MAP direction with 0.998 cosine similarity under realistic finite-shot noise.

The open question connecting both lines of work is where any of this should actually run. A QPU being available is not the same as it being worth using once queueing, compilation, shots, and mitigation are counted honestly. I am currently building a resource-aware execution policy, grounded in a systematic review of the quantum-HPC integration literature, that weighs a QPU route's complete cost against classical and simulated alternatives and is explicitly allowed to say no. Across all three lines, the goal is the same: not to make more things quantum, but to know, with evidence, exactly when it is worth it.

Education

The University of Texas at San Antonio (UTSA) — *Aug. 2023–Present*

Ph.D. Student, Electrical Engineering — *San Antonio, TX, USA*

- Advisor: Prof. Panagiotis P. Markopoulos.
- Dissertation: *Structure-Aware and Resource-Efficient Quantum–Classical Learning: Multilinear Architectures, Adaptive Measurement, and Resource-Aware Execution.*
- GPA: 3.93/4.00.

Vietnam National University HCMC–University of Information Technology — *Aug. 2018–Jan. 2022*

B.S. (Honors), Computer Science — *Ho Chi Minh City, Vietnam*

- Graduated with Excellence; ranked 1st of 26 students; final grade: 9.05/10.
- Thesis: *Evaluating Deep Learning Methods for Vietnamese Infographic Visual Question Answering.*
- Advisors: Tien Do Van, M.Sc., and Thanh Duc Ngo, Ph.D.

Publications and Manuscripts

Accepted / To Appear

1. **Van Tien Nguyen** and Panagiotis P. Markopoulos, "VQBR: Variational Quantum Bayesian Regression via Measurement-Based MAP Direction Recovery," in *Proceedings of the 2026 IEEE International Conference on*

Quantum Computing and Engineering (QCE), Quantum Machine Learning Technical Papers Track, Toronto, Canada, Sept. 2026. *Accepted.*

2. **Van Tien Nguyen** and Panagiotis P. Markopoulos, "Parameter-Efficient Multilinear Readout for Quantum Reservoir Computing," *Workshop on Quantum Machine Learning at IEEE Quantum Week 2026 (QML@QCE)*, Toronto, Canada, Sept. 2026. *Accepted for speed presentation and poster; to appear in workshop proceedings.*
3. **Van Tien Nguyen**, Mayur Dhanaraj, and Panagiotis P. Markopoulos, "Depth- and Rank-Selective Multilinear Transformation Layers for Compact CNNs," *60th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, Oct. 2026. *Accepted.*

Peer-Reviewed Publications

4. **Van Tien Nguyen** and Panagiotis P. Markopoulos, "Quantum-Based Tensor Contraction Layers," in *Machine Learning from Challenging Data 2026*, Proc. SPIE, vol. 14030, paper 140300I, 2026. doi:10.1117/12.3098024.
5. M. Tran, T. Pham, **V. T. Nguyen**, T. Do, and T. N. Duc, "A Robust Framework for Mathematical Formula Detection," in *2021 International Conference on Multimedia Analysis and Pattern Recognition (MAPR)*, pp. 1–6, 2021. doi:10.1109/MAPR53640.2021.9585197.

Preprints and Manuscripts in Revision

6. **Van Tien Nguyen** and Panagiotis P. Markopoulos, "Quantum Tensor Decomposition: A Review of Foundations, Algorithms, Challenges, and Future Directions," manuscript in preparation, 2026.
7. Mayur Dhanaraj*, **Van Tien Nguyen***, and Panagiotis P. Markopoulos, "Adaptive Low-Rank Multilinear Transformations for Compact Convolutional Neural Networks", manuscript in revision for resubmission to *IEEE Transactions on Artificial Intelligence*, 2026. **Equal contribution; supported by NSF Award 2332744.*
8. **Van Tien Nguyen** and Panagiotis P. Markopoulos, "Quantum Optimization for Adaptive Low-Rank CP Decomposition," manuscript in revision, 2026.

Research Experience

Graduate Research Assistant, Machine Learning Optimization and Systems (MILOS) Lab — *Aug. 2023–Present*

The University of Texas at San Antonio — *San Antonio, TX*

- Develop tensor-structured and measurement-efficient methods for quantum and hybrid quantum–classical learning under limited qubits, finite shots, and noisy hardware.
- Designed VQBR, which separates a Bayesian regression coefficient vector into direction and scale, estimates a Gram-matrix-free objective through overlap and Pauli measurements, and recovers scale analytically; validated the method on simulators and IBM Quantum hardware.
- Developed Tucker-factored readouts for quantum reservoir computing, reducing trainable parameters by approximately 12.1× relative to an MLP readout and achieving compression ratios up to 52.24× across tested reservoir sizes.
- Co-developed depth- and rank-selective multilinear layers for compact CNNs and evaluated parameter, computation, and predictive-performance tradeoffs on CIFAR-10, DOTA-v1.0, and ImageNet-1K.
- Formulated interaction-aware adaptive tensor-rank selection as a QUBO problem for structured CP decomposition, benchmarked against exhaustive search and classical heuristics.
- Leading a literature-grounded synthesis of quantum-HPC integration, middleware, and workflow-orchestration systems to motivate a learning-aware execution policy for routing subtasks across CPU, GPU, simulator, and QPU resources under complete end-to-end cost.
- Investigate quantum-assisted CP/PARAFAC decomposition and objective-aligned ambiguity resolution for recovering deployable classical predictors from quantum-encoded states.
- Build reproducible workflows in PyTorch, Qiskit, and PennyLane; support shared GPU/CPU infrastructure, software environments, and remote experimentation across the laboratory.

Research Assistant, AI Club — *2020–2022*

VNUHCMC–University of Information Technology — *Ho Chi Minh City, Vietnam*

- Conducted deep-learning research in computer vision and NLP, including infographic visual question answering, instance retrieval, mathematical-formula detection, and reinforcement-learning game agents.
- Contributed to dataset construction, model training, evaluation, technical writing, and research presentations.

Industry Experience

Research Engineer — *Mar. 2022–Mar. 2023*

Viettel High Technology Industries Corporation — *Vietnam*

- Developed NLP systems for named-entity recognition, entity linking, summarization, and timeline generation using internally constructed news and social-media datasets.
- Packaged models as containerized microservices and integrated them into production workflows using version control,

CI/CD, and reproducible model-serving practices.

Teaching Experience

Incoming Graduate Teaching Assistant, Department of Electrical Engineering *Aug. 2026–May 2027*
The University of Texas at San Antonio *San Antonio, TX*

- Appointed for the 2026–2027 academic year; assigned to EE 2423 Network Theory for Fall 2026.
- Responsibilities include grading, office hours, student support, course logistics, and preparation of instructional materials.

Teaching Assistant *Sep. 2021–Jan. 2022*
VNUHCMC–University of Information Technology *Ho Chi Minh City, Vietnam*

- Supported tutorials, grading, office hours, and student programming exercises for Introduction to Programming in C++ and Python.

Patents and Inventions

- **Quantum-to-Classical Ambiguity Resolution (QCAR)**, UTSA invention disclosure, 2026. Objective-aligned recovery of missing sign and scale information from quantum-encoded model parameters to produce an explicit classical predictor.
- **Detecting Ships in Distress Based on Optical Satellite Images and Automatic Identification System Signals**, Vietnamese patent, Jan. 2023.

Honors and Awards

- **NSF Student Travel Grant, IEEE Quantum Week 2026**, up to $1,250 in reimbursable travel support to attend and present accepted work at the IEEE International Conference on Quantum Computing and Engineering (QCE 2026), Toronto, Canada, July 2026.
- **Selected for Phase III**, qBraid–MITRE–JonesTrading Quantum Challenge / Global Industry Challenge 2026, June 2026.
- **UT San Antonio Institutional Nominee, 2026 Google PhD Fellowship–Quantum Computing**, selected as one of four university nominees, Apr. 2026.
- **Quantum Rising Star Award**, inaugural Quantum Student Success Summit (QS3), Rice University, Mar. 2026.
- **Graduate School Academic Travel Fund**, UTSA, $750 to support presentation at SPIE Machine Learning from Challenging Data 2026, Spring 2026.
- **Earlier distinctions:** Student of 5 Merits, University of Information Technology, 2019–2021 (three consecutive years); Consolation Prize, AI Camera Inspection Challenge, 2020; Consolation Prize, National Mathematical Olympiad Competition, 2019.

Professional Leadership, Certification, and Service

- **Qiskit Advocate**, IBM Quantum / Qiskit community, Aug. 2025–Present; advanced to Tier 1 in July 2026.
- **Organizer**, Qiskit Fall Fest, The University of Texas at San Antonio, 2026.
- **IBM Certified Quantum Computation using Qiskit v2.X Developer–Associate**, June 2026. Credential.
- **Mentor-in-Training**, Qiskit Global Summer School 2026, IBM Quantum.
- Presenter, SPIE Defense + Commercial Sensing, Machine Learning from Challenging Data, National Harbor, MD, Apr. 2026.
- Poster presenter and lightning-talk speaker, inaugural Quantum Student Success Summit, Rice University, Mar. 2026.
- Participant, Qiskit Global Summer School, 2024 and 2025; earned Quantum Excellence completion badges.
- Participant, NSF Spring AI School, San Antonio, TX, 2024–2026.
- Scribe, EXAIL Workshop, Pittsburgh, PA, Oct. 2024.

Selected Additional Technical Projects

- **Transformer-Based Quantum Error Mitigation:** Benchmarked unmitigated estimation, zero-noise extrapolation, a histogram MLP, and a Transformer encoder on raw shot sequences; a shuffle ablation showed that the Transformer learned a permutation-invariant representation. Code and report.
- **NV-Center Quantum Sensing for Biomedicine:** Simulated nanodiamond NV-center T_1 relaxometry for mitochondrial free-radical conditions using amplitude damping, inversion recovery, T_1 fitting, and standard-quantum-limit ensemble scaling. Code and report.
- **Quantum Neural Networks Benchmark Study:** Implemented variational quantum classifiers and compared them with classical baselines across multiple benchmark datasets. Code and report.
- **HHL for Linear Regression:** Implemented and analyzed HHL-based linear regression on quantum simulators, including numerical stability, conditioning, data-loading assumptions, and resource requirements. Code and report.

Technical Skills

Programming / ML	Python, C/C++, PyTorch, TensorFlow, scikit-learn, NumPy, SciPy, snnTorch
Quantum Computing	Qiskit, PennyLane, Cirq, CUDA-Q, IBM Quantum Runtime; variational algorithms and quantum simulation
Research Computing	Linux, Git, GitHub, Docker, CUDA/GPU workflows, Jupyter, LaTeX, CI/CD