

Tushar Pandey, PhD

Dallas, TX | wtatushar@gmail.com | linkedin.com/in/tpmath | pandey-tushar.com | scholar.google.com/citations?user=_ivaRglAAAAJ

EDUCATION

Texas A&M University, Ph.D., Mathematics (Advisor: Prof. Tian Yang) | Aug 2024

Indian Institute of Technology (IIT) Kanpur, B.S., Mathematics and Scientific Computing | Jul 2019

- **Relevant graduate work & TA:** Mathematical Foundations for Data Science (MATH 677 TA), Stochastic Processes, Convex Optimization, Topological Data Analysis, Advanced Calculus, Probability.

PUBLICATIONS

- *Adaptive Graph of Thoughts: Test-Time Adaptive Reasoning Unifying Chain, Tree, and Graph Structures*, arXiv:2502.05078 (2025). **First author.** Test-time adaptive reasoning dynamically composing chain/tree/graph structures from intermediate-step signal; +400% benchmark accuracy; **19 citations**; NeurIPS / ICLR submission in preparation.
- *ML-Driven Quantum Hacking of CHSH-Based QKD: Exploiting Entropy Vulnerabilities in Self-Testing Protocols*, QUEST-IS'25 (Springer). **First author.** Adversarial ML against entropy assumptions in cryptographic protocols.
- *Modeling Sea Level Rise Using Ensemble Techniques*, Analytics 3(3):276–296 (2024). Time-series ensemble Machine Learning on million-row datasets; 1st prize at TAMIDS data-science competition.
- *Quantum Kernels for Parity-Structured Classification*, IEEE DCAS (2026). *Fixed-Reservoir vs. Variational Quantum Architectures for Chaotic Dynamics*, arXiv:2604.23743 (2026). *Fair Decoder Baselines for Bivariate Bicycle Codes*, arXiv:2603.19062 (2026).
- *Topological Phases of Matter, Quantum Error Correction, and Topological Twists*, APL Quantum 2(3):031501 (2025). *Efficient VQE Methodologies on Quantum Processors*, arXiv:2407.16107 (2024); plus 4 additional first/sole-author publications in mathematical reasoning and quantum information.
- **Surveys:** *Quantum Computing through ChatGPT* (2023); *Topological Understanding of Neural Networks* (2022).

INDUSTRY & RESEARCH EXPERIENCE

Senior Software Engineer, DataRobot (Software Developer II, Mar 2025 – Mar 2026) | Mar 2026 – Present

- Architecting **multi-tenant Large Language Model (LLM) inference at scale**, designing inference-quota and serving-capacity systems that share GPU/CPU pools across users for **real-time, efficient model serving** without head-of-line blocking.
- Reduced model deployment latency **30–40%** across enterprise workloads via a distributed inference-caching layer; delivered **300% throughput uplift** by replacing VM-based serving with a container-native alternative; security isolation maintained.
- Drove production certification of **Covalent** for end-to-end orchestration of Generative AI pipelines; integrated customer-driven model requirements into the platform; serving as Subject-Matter Expert for high-impact platform services.
- Automated cross-repository pipelines for onboarding, CI/CD, and CVE remediation; resolved 1,500+ critical CVEs; produced internal tooling adopted across engineering teams.

Research Software Engineer, Agnostiq (Covalent distributed compute) | Jul 2024 – Feb 2025

- Authored and shipped **Adaptive Graph of Thoughts (AGoT)** end-to-end, from research idea to production agentic AI blueprint; the framework **co-designs the reasoning algorithm and the inference-time interface** to react in real-time to product signal.
- Cut LLM **LoRA fine-tuning cost 99%** via CPU-first federated-learning architecture (matched quality at <1% original spend), model efficiency at training time.
- Shipped **10+ production agentic AI blueprints**: Retrieval-Augmented Generation (RAG), multi-agent planning, tool-use orchestration, distillation pipelines, optimized inference, reasoning-over-documents, reusable patterns adopted across enterprise customers.

Summer Researcher (NSF–MSGI Fellow), Oak Ridge National Laboratory (ORNL) | Summers 2021 & 2022

- Built quantum-ML / topological-data-analysis pipelines for ground-state Hamiltonian classification; developed mathematical framework for topological phase transitions applicable to ML pattern recognition (APL Quantum 2025, lead author).

Graduate Researcher & TA (Ph.D., Mathematics), Texas A&M University, Advisor: Tian Yang | Aug 2019 – Jul 2024

- 5 years of independent research producing first-author publications in reasoning, asymptotic analysis, and optimization; founded directed-reading program mentoring 20+ students in Machine Learning, convex optimization, computational topology, and TDA.

SERVICE, TALKS & OPEN-SOURCE PROJECTS

- **Organizer**, AMS Fall Eastern Sectional 2025 Special Session on QEC / Categorical Approaches to Topological Codes. **Invited mentor**, NYUAD Quantum Hackathon 2025. **Invited talks**: JMM 2024, IBM Quantum Developer Conference 2025, Yale, UCSB, ICTP (Italy), ICERM (Brown), ICMAT (Madrid), Rutgers, Michigan State, AMS Sectionals.
- **Open-source projects**: AGoT (github.com/AgnoStiqHQ/multi-agent-llm); NeurotiQ, Quantum SVM for EEG; Alzheimer Detection via deep learning on biomarkers; Harmful Brain Activity classification; QMPRS, linear-depth MPS encoding (QHACK / Zenodo); Sea Level Rise Streamlit app; ZX-Live (Quantomatic open-source); Compressed Sensing; Approximating Credit-Guards; Quantum ODE solver for large linear systems (2026).

AWARDS & RECOGNITION

- **IBM Quantum Developer Conference 2025** Challenge Winner (Schwinger-model hadron dynamics); **Q-Volution Hackathon 2026** 1st Prize (GiQ); **MIT iQuHack 2024** 1st Prize (Classiq general-arithmetic challenge); Top-15 QHACK Quantum Algorithms Challenge (2023).
- **TAMIDS Data Science Competition**, 1st Prize 2021 (+Best Presentation), 2nd Prize 2022 (+Best Presentation), 3rd Prize 2024; TAMU Datathon 1st Prize HPE Challenge (2020); **Guseman Prize**, TAMU Math Dept (2023); TAMU Graduate Researcher Award (2023); Special Appreciation Award, IIT Kanpur (2019).

SKILLS

AI / Machine Learning: Large Language Models, Generative AI, Agentic AI, Test-Time Adaptation & Reasoning, Retrieval-Augmented Generation (RAG), LoRA fine-tuning, Federated Learning, Distillation, Model Efficiency (training & inference), Prompt Engineering, At-Scale Model Serving, Multi-Tenant Orchestration, PyTorch, TensorFlow, scikit-learn, Transformers, NLP, ensemble methods, time-series modeling, TDA, statistical learning

Programming: Python (primary, NumPy, pandas, SciPy, statsmodels), Go, Java, SQL, MATLAB; bash, Linux

Platform / MLOps: Docker, Kubernetes, CI/CD, AWS / Azure / GCP, Covalent distributed compute, Streamlit, Git, GitHub Actions, Jupyter, Hugging Face