Distributed Classical–Quantum Algorithms for Certifiable Trustworthy Machine Learning

Overview

Modern machine learning systems are increasingly used in high-stakes decision-making, yet many learning algorithms remain heuristic, opaque, and difficult to certify. Interpretable models such as decision trees, clustering models, and decision-tree-based policies offer greater transparency, but training these models to global optimality is computationally challenging because the underlying problems are often large-scale, nonconvex, and combinatorial. My long term goal is to develop a new algorithmic foundation for certifiable machine learning through distributed classical–quantum optimization. The central goal is to enable trustworthy and interpretable learning models to be trained with provable feasibility, optimality, or near-optimality certificates, rather than relying only on heuristic search. The research will be organized around four integrated thrusts. i) develop reduced-space formulations and certificate-generating algorithms for interpretable learning problems, including optimal decision trees, constrained clustering, and interpretable reinforcement learning policies. ii) design distributed decomposition methods that split large learning problems into smaller subproblems while maintaining global consistency and certificate tracking. iii) create quantum-assisted and quantum-inspired optimization modules for selected decomposed subproblems. iv) integrate neural guidance with certificate-preserving global optimization to accelerate search without sacrificing reliability.

The education plan is tightly integrated with the research vision. The project will create a new educational pathway on certifiable and quantum-enabled trustworthy AI. This will include new course modules in deep learning analytics and optimization-based machine learning, undergraduate research projects on interpretable and certifiable AI, and open-source teaching materials that visualize branch-and-bound, QUBO reformulations, quantum-assisted optimization, and optimality-gap certificates. The project will also establish a student research pipeline that trains undergraduate and graduate students to combine machine learning, optimization, quantum computing, and trustworthy AI. Through classroom integration, mentoring, open-source software, and outreach activities, the project will broaden student participation in emerging areas at the intersection of AI and future computing.

Intellectual Merit

The intellectual merit of this project lies in establishing a principled computational framework for training interpretable machine learning models with certificates of quality at scales that are beyond the reach of traditional exact optimization methods. The project will contribute new theory, algorithms, and computational architectures for certifiable learning.

First, the project will develop a unified reduced-space global optimization theory for interpretable machine learning. Many learning problems, including optimal decision tree training, constrained clustering, and interpretable policy learning, can be represented as structured two-stage optimization problems. The proposed work will characterize when these formulations admit compact structural search spaces, decomposable lower bounds, efficient upper bounds, and convergent reduced-space branch-and-bound algorithms. These results will provide new theoretical foundations for exact and near-exact learning algorithms.

Second, the project will create scalable distributed decomposition methods for large-scale machine learning optimization. Instead of treating a learning problem as one large mixed-integer formulation, the project will develop sample-level, scenario-level, tree-structure-level, and cluster-level decomposition strategies. These methods will enable parallel computation of bounds, adaptive sample aggregation, structure-aware partitioning, and feasibility-preserving solution merging. The resulting algorithms will expand the practical scale of certifiable learning while maintaining interpretable model structures and rigorous optimality-gap information.

Third, the project will advance distributed classical–quantum optimization for machine learning. Current quantum devices are not large enough to solve full-scale learning problems directly. This project addresses that limitation by developing decomposition methods that generate many small QUBO or Ising subproblems compatible with available quantum and quantum-inspired hardware. Quantum annealing, QAOA, coherent Ising machines, and quantum-inspired solvers will be studied as acceleration modules within a larger certificate-preserving optimization framework. This approach will move quantum optimization beyond isolated toy instances toward distributed, verifiable machine learning pipelines.

Fourth, the project will develop neural-guided but certifiable optimization methods. Neural networks, graph neural networks, transformers, or large-language-model-inspired modules may provide effective guidance for branching, decomposition, warm starts, and variable fixing. However, purely learned solvers often lack reliability guarantees. This project will design mechanisms by which learned predictions accelerate the search while exact optimization modules verify feasibility, repair invalid solutions, and maintain optimality certificates. This creates a new paradigm of learning-guided optimization with certification, combining the speed of data-driven search with the rigor of mathematical optimization.

Collectively, these contributions will generate new knowledge at the intersection of algorithm design, global optimization, machine learning, quantum computing, and trustworthy AI.

Broader Impacts

The broader impacts of this project are threefold: advancing trustworthy AI, preparing a future computing workforce, and broadening access to interdisciplinary research opportunities.

First, the project will support the development of trustworthy AI systems by creating algorithms that are both interpretable and certifiable. In high-stakes domains such as healthcare, transportation, manufacturing, energy, and public-sector decision-making, users need models that can be understood, audited, and trusted. The proposed methods will enable interpretable models such as decision trees, clustering models, and policy trees to be trained with explicit certificates of feasibility and solution quality. This can reduce reliance on opaque black-box systems and support more transparent decision-making.

Second, the project will train students in an emerging interdisciplinary area that combines machine learning, optimization, quantum computing, and trustworthy AI. New course modules will introduce students to certifiable machine learning, reduced-space global optimization, QUBO and Ising reformulations, quantum-assisted optimization, and neural-guided search. These materials will be incorporated into existing courses and made available as open educational resources. Students will learn not only how to build AI models, but also how to evaluate their reliability, interpretability, and computational guarantees.

Third, the project will create hands-on research opportunities for undergraduate and graduate students. Undergraduate students will participate in accessible research projects such as visualizing branch-and-bound, comparing heuristic and globally optimized decision trees, implementing small QUBO models, and studying quantum-inspired optimization methods. Graduate students will contribute to algorithm design, theoretical analysis, software development, and computational experiments. This mentoring structure will help build a diverse pipeline of students prepared for careers in AI, optimization, data science, quantum computing, and advanced computing systems.

The project will also produce open-source software and educational notebooks that allow students, researchers, and practitioners to experiment with certifiable learning algorithms. These resources will include demonstrations of optimal decision trees, constrained clustering, interpretable policy learning, QUBO reformulations, quantum-assisted subproblem solving, and optimality-gap visualization. By making these tools publicly available, the project will lower the barrier for learning and research in certifiable AI.

In the long term, this CAREER project will help shift the practice of machine learning from heuristic model training toward certificate-aware, interpretable, and computationally reliable AI. It will also prepare students to lead in a future where AI, optimization, and quantum-enabled computing are increasingly interconnected.