

Q-NEST: Quantum Neuroevolutionary Strategies for Parametrically Optimizing and Topologically Augmenting Artificial Neural Networks via Novel Mutation, Translocation, and Crossover Hybrid Quantum-Classical Genetic Operators

Tara, Paarth (School: North Carolina School of Science and Mathematics)

In the pursuit of increasingly robust artificial intelligence, it is useful to look towards nature for inspiration. Neuroevolution does just this by simulating Darwinian-style natural selection in populations of artificial neural networks to effectively yield global optima. Compared to traditional stochastic gradient descent (SGD) methods, neuroevolution has proven to more rapidly converge to global solutions, especially within deceptive, sparse fitness landscapes. However, the computational requirements of neuroevolution make it impractical to implement on classical machines. Quantum parallelization via entanglement, superposition, and phase-interference offers a unique solution to this constraint by offering a robust platform to perform neuroevolutionary computations. I present Q-NEST (Quantum Neuroevolutionary Strategies), a hybrid quantum-classical framework capable of rapidly optimizing populations of quantum neural networks (QNNs) under an arbitrary objective function via novel mutation, translocation, and crossover operators. Q-NEST utilizes repeat-until-success quantum neurons, constructions of artificial quantum chiasmata, perturbative mutation circuits, and techniques in topological augmentation and speciation to efficiently converge to solutions, especially within noisy cost functions littered with local optima. Numerical experiments demonstrate that Q-NEST outperforms random search to effectively implement neuroevolutionary optimization and rapidly converge to global optima in both noisy and noiseless quantum architectures. Q-NEST serves as an exceptionally promising direction in the realm of quantum machine learning by revealing the potential of neuroevolution on variational quantum circuits that mimic non-linear, dissipative neural dynamics.

Awards Won:

