Day 3

Kyle Mills, 1QBit, Vector Institute for Artificial Intelligence

Title: Finding the ground state of spin Hamiltonians with reinforcement learning

Abstract:

We demonstrate the use of reinforcement learning (RL) in solving optimal control problems for the evolution of classical and quantum systems. We apply deep RL to simulated annealing (SA), demonstrating that a RL agent can, through experience alone, devise a temperature schedule for finding ground states of a model, surpassing the performance of standard heuristic temperature schedules.

We demonstrate that the RL agent can be trained solely from interacting with a limited number of spin glass instances. The RL agent is then used to drive the temperature of a Monte-Carlo method on test instances never previously shown to the agent. Additionally, the resulting algorithm scales (with respect to system size) much more favourably.

We demonstrate the robustness of the RL approach when the system operates in a “destructive observation” mode, alluding to a quantum system where measurements destroy the state of the system. Analogous quantum control problems in adiabatic quantum computing (e.g. controlling the transverse field in quantum annealing, or scheduling gate rotations in QAOA) could be tackled with such an approach, suggesting a promising research direction combining reinforcement learning with adiabatic quantum computing.

The results of this investigation have been published in Nature Machine Intelligence 2, 509-517 (2020).