

Variational neural annealing

Many combinatorial optimization problems relevant to computer science, computational biology and physics can be tackled with simulated annealing, which is a powerful framework for optimizing the properties of complex systems through the lens of statistical mechanics. However, simulated annealing and its quantum counterpart, simulated quantum annealing, are traditionally implemented via Markov chain Monte Carlo, often displaying slow convergence to optimal solutions for challenging optimization problems. Here we combine the variational principle in classical and quantum physics with recurrent neural networks, whose dynamics are naturally devoid of slow Markov chains to accurately emulate annealing in its classical and quantum formulations. We find that a variational implementation of classical annealing is not only superior to its quantum analog in terms of speed of convergence and accuracy of solutions but also outperforms traditional simulated annealing and simulated quantum annealing on several prototypical spin glass models as the annealing time is increased. In addition, through the use of the sign-problem-free variational Monte Carlo method, our framework can emulate quantum annealing with non-stoquastic drivers at moderately large system sizes, thus providing a useful tool to benchmark the next generation of quantum annealing devices which implement non-stoquastic Hamiltonians.