Benchmarking Variational Quantum Algorithms for Combinatorial Optimization in Practice

T Schwägerl, Y Chai, T Hartung, K Jansen, S Kühn

Variational quantum algorithms (VQAs) and, in particular, variants of the varational quantum eigensolver (VQE) have been proposed to address combinatorial optimization problems. Using only shallow ansatz circuits, these approaches are deemed suitable for current noisy, intermediate-scale quantum hardware. However, the resources required for training shallow variational quantum circuits often scale superpolynomially in problem size, due to the presence of many poor local minima. In this study we numerically investigate what this scaling result means in practice for solving combinatorial optimization problems using MaxCut on 3-regular graphs as a benchmark. For fixed resources, we compare the average performance of training a shallow variational quantum circuit, sampling with replacement and a greedy algorithm starting from the same initial point as the quantum algorithm. We identify a minimum problem size for which the quantum algorithm can consistently outperform sampling and, for each problem size, characterize the separation between the quantum algorithm and the greedy algorithm. Furthermore, we extend the average case analysis by investigating the correlation between the performance of the algorithms by instance. Our results provide a step towards meaningfully benchmarking the performance of variational quantum algorithms for combinatorial optimization problems for a realistic set of resources.

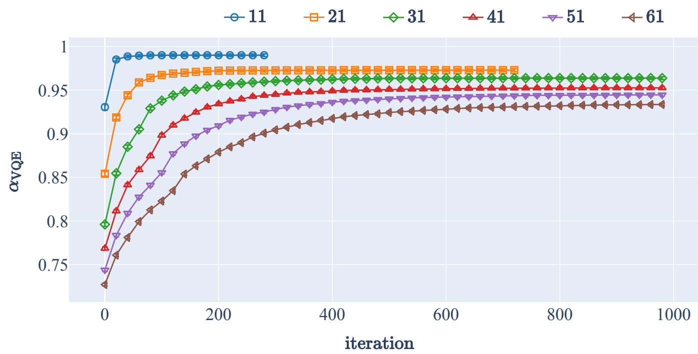

(a) Average approximation ratio α_{VQE} achieved by the VQE as a function of its iteration count. The different markers encode different problems sizes. High approximation ratios, such as for size 11, can be misleading because they are in fact worse than randomly sampling solutions.

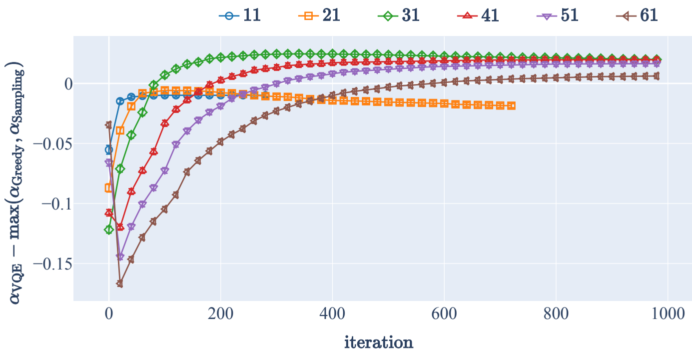

(b) Average difference in approximation ratio of the VQE and the best one obtained by sampling and a greedy algorithm. Now, it is visible that the VQE performs worse than the competing algorithms for small problems. Only for larger problems it can consistently outperform sampling and the greedy algorithm.

Figure 1: Summary of our main result. The approximation ratio α is frequently used to study the performance of VQAs for combinatorial problems, such as MaxCut on 3-regular graphs. However, it can be a misleading metric if one does not carefully take into account the resources used by the algorithm. We perform a more meaningful benchmark by comparing the VQE to sampling with replacement and a greedy algorithm starting from the same initial point as the VQE for fixed resources.