Separating quantum and classical computing: rigorous proof and practical application

Marshall, S.C.

Summary

This thesis has sought to probe or prove the supposed advantages of quantum computing, first by analysing under what conditions an advantage may theoretically be present, and then by developing and scrutinising techniques to produce a practical advantage.

The introductory chapter introduces the core concepts needed to understand what quantum computing is, how it fits into concepts like complexity theory or machine learning, as well as more developed quantum computing topics, such as circuit cutting.

The second chapter develops the strongest evidence given so far that quantum computers can sample from distributions or compute functions that classical computers cannot, which we formalise as $\mathsf{SampP} \stackrel{?}{=} \mathsf{SampBQP}$ or $\mathsf{FBQP} \stackrel{?}{=} \mathsf{FBPP}$. This chapter iterates on existing conditions for the separation of quantum and classical computation, finding stronger conditions than had previously been achieved and suggesting that if quantum computers cannot provide speed-ups over classical, then widely held beliefs about complexity theory would be wrong. This progress is made by demonstrating that exact counting and approximate counting have to be different unless the polynomial hierarchy collapses to its second level.

Chapter 3 follows the lead of Chapter 2 by questioning if advice may change the separation of classical and quantum computation. Motivated by papers that identify tasks which have a likely quantum advantage only need a quantum advantage to prepare samples, such as [38]. This chapter first generalised this question, to an advice generating Turing machine

and an advice recieving machine. This generalisation allowed for general rules, particularly $\mathsf{P}/\mathsf{poly}^{\mathsf{B}} = \mathsf{P}^{\mathsf{Un}(\mathsf{B})}$. This result reveal that advice from quantum computers enhances classical computers, unless $\mathsf{BQEXP}{=}\mathsf{BPEXP}$ (if quantum and classical machines are equally useful if they are both allowed to run for exponentially long). This chapter highlighted the practical implications of quantum advice givers, as quantum computers will initially be very expensive being able to use them for a small portion of the total compute of some quantum-enhanced task will be key.

Chapters 4 and 6 can be seen as building on the practical use case of the previous chapter, by both attempting to enhance limited modern machines. Chapter 6 does this directly by developing a model which uses a quantum machine to prepare/train a model that can then be deployed using only classical hardware. The model is shown to be universal for models prepared quantumly and deployed classically, therefore the results of Chapter 3 imply it is more powerful than classical models unless $\mathsf{BQEXP}{=}\mathsf{BPEXP}$. Chapter 4 enhances limited modern machines via circuit cutting, where a model is developed that in the limit becomes the cut-up version of a parameterised quantum circuit but can also be set to use much less resources than perfect circuit cutting would require. Through experiment, it is demonstrated that this model is capable of learning distributions which are non-trivial for a quantum machine (such as handwriting), while still retaining advantages from the quantum machine (such as being able to fit quantum circuits). Both of these chapters probe the power of quantum computing by opening up the ability for us to study quantum algorithms that we would otherwise not be able to study on modern machines.

Chapter 5 takes a different tone than the previous chapters: Instead of furthering a research line to show that quantum computing is more powerful than classical computing, it suggests that a branch of research, circuit cutting, can never provide a meaningful speed up. This is achieved by showing that if circuit cutting could be done efficiently, then $\mathsf{BQP} = \mathsf{BPP}$. This has direct implications for the results of Chapter 4, showing this model can never be as good as a full-sized model for all tasks. It also helps to bound and inform a number of improvements to cutting schemes.

All of these chapters either demonstrate quantum-classical separations or provide results informative to those trying to demonstrate such a connection.