Srinivasan Arunachalam

Given a circuit, we can always treat it as a black box which output a quantum state with given inputs. If the circuit is promised to be in a certain class, (for instance threshold circuits, AC 0 circuits or IQP circuits), learning the structure of the circuit is believed hard classically (quasi polynomial time, $n^{\log(n)}$). This is possibly due to the reason of difficulty of factoring, Fourier sampling and etc. It would be interesting to see if this can be learned efficiently with quantum computers. If this is the case, this could mean that we can learn neural networks faster quantumly, since the circuits can always be treated as a neural network with different parameters specifying different gates.

It would also be interested to search further for quantum inspired classical algorithm.

Fernando Brandão

An interesting recent result is the work by S. Bravyi et al. (arXiv:1704.00690), where the authors introduced the 2D Hidden Linear Function problem that could be solved using constant depth quantum circuits whereas classical circuits must have logarithmic depth. This could have implications of the separation between quantum neural networks and classical neural networks. For instance, one could possibly come up with a dataset, which can be represented using constant depth of quantum circuits whereas the classical circuits would cost logarithmic depth. If phrased with generative models, we could also look for different outputs that can be generated by quantum circuits and classical circuits, which may lead to a separation.

Another interesting problem is that, with the support vector machine, the kernel matrix is often used as a crucial element to implement the optimization. It has been shown that once the kernel matrix is prepared, exponential speedup can be achieved with quantum computers in the optimization (*Phys. Rev. Lett.* ***113****, 130503*). The problem then turns into how to prepare the kernel matrix or even how to define a suitable kernel matrix given the dataset. With a quantum computer, the kernel could be prepared by the Gibbs state of some local Hamiltonian. It would be interesting to see if there is any relevant machine learning problem that could be done with the NISQ devices.

Rolando Somma

There has been some progress on the use of machine learning techniques for problems in quantum information. Two interesting examples are shorter-depth quantum algorithms for certain problems and better ansatz for eigenstates of Hamiltonians. The use of quantum computers to speedup calculations in machine learning is still an open problem.

It would also be interesting to understand the importance of quantum algorithms for linear algebra problems (e.g., HHL and related) in various scientific disciplines.

Scott Aaronson

One of the most important advances is the breakthrough of Ewin Tang's (arXiv:1807.04271) dequantization of Kerenidis-Prakash algorithm for recommendation system (arXiv:1603.08675). The algorithm recommends products to users based on a matrix whose rows are lists of the users and columns are the products. The crucial assumption for the algorithm to work is that this matrix is close to some low rank matrix. If one could only randomly access the matrix elements, nothing better could be done than a Grover type search (arXiv:quant-ph/9701001). If one wants a speedup beyond what Grover search can do, additional assumptions must be made about how the data could be accessed. Kerenidis-Prakash algorithm assumes that the know entries of the matrix are sorted into a binary search tree data structure. However this assumption is still controversial. Kerenidis-Prakash algorithm is able to sample recommendations for given users in logarithmic time.

The work of Tang is thought to be the first classical end to end application, unlike the HHL where the solution resides in the amplitude and still needs to be measured. It uses the same data structure as the Kerenidis-Prakash algorithm. Using the property that the singular values of a big low rank matrix could be well approximated by its submatrix's singular values (JACM, 51, 1025–1041, 2004), Tang's algorithm is able to make query for given users in polylogarithmic time. Thus, there is still a polynomial gap between quantum and classical running times.

This is the kind of work that needs to be done for quantum machine learning, which is the end to end applications of the quantum machine learning primitives. To put it in another way, is there an application, which starts with classical inputs, and after you run the quantum algorithm, one actually get some output that people care in the classical world. One might even prove that in such case, no classical algorithm can solve the problem efficiently.

The field is wide open to try to look for examples that for some quantum machine learning algorithms, we can prove the lower bounds of classical randomized query complexities, which might leads to a provable exponential separation.

Eleanor G. Rieffel

It is really encouraging to see the results proving that classical computers cannot sample efficiently from even quite simple quantum circuits (*Phys. Rev. Lett.* ***117****, 080501*), particularly the robustness results for IQP circuits. It would be great to see experiments with emerging hardware that takes us beyond what can be simulated classically.

In terms of important open problems for the near term, some of the exciting questions are, what are the best experiments we could perform on NISQ devices to get insight into the most promising long term directions for quantum machine learning, what will be the earliest application of quantum machine learning, what resources will it require (qubits, quantum operations, hardware architecture, robustness, etc.), and how effectively can we use classical machine learning to set parameters in (or more generally design) quantum algorithms?

Seth Lloyd

With the kernel methods, the data is embedded in a high dimensional Hilbert space. Then, we can use the intrinsic methods of this Hilbert space to try to find characteristics about the data. Problems that are nonlinear if we only have low dimensional representations of the data become linear if we embed the data in a high dimensional Hilbert space. This makes it possible for quantum computers to solve problems that are dimension dependent, for instance clustering problems, which we cannot really do with classical computers.

The most enabling development that could take place in quantum machine learning would be the experimental construction of medium to large scale qRAM. For instance, if we have thousands of training examples, it would require roughly ten qubits, which would not be necessarily a very large system and is possibly able to be built. In principal, qRAM does not require error correction for reasonable amount of data, where acceptable error tolerance would be around 1/log(N) (N as the number of data points).

What is role of quantum annealer for quantum machine learning tasks? Even with classical deep learning, it was not working very well for years without the computational power we are having nowadays. Thus, we should expect progresses once we had deep quantum learners, which is possible in the near future.

Nathan Wiebe

Can quantum computers provide a super-polynomial speedup for a practical problem of interest to the classical machine learning community?

What is the expressive power of current proposals for quantum neural networks relative classical neural networks on classical rather than quantum data?

What are the overheads at an experimental or fault tolerant level facing the construction of a medium-scale QRAM capable of storing a few thousand training examples?

Anupam Prakash

The most important advances over the past year would be the development of quantum algorithms for convex optimization.

It would be interesting to see natural optimization problems where quantum algorithms can achieve a provable polynomial speedup.