

# 41076: Methods in Quantum Computing

Maria Kieferova

July 2020

## 1 Hamiltonian simulation

Given an initial state of a system  $|\psi(0)\rangle$  and a Hamiltonian  $H$ , our goal is to simulate the time evolution  $|\psi(0)\rangle \rightarrow e^{-iHt} |\psi(0)\rangle$ . The goal of Hamiltonian simulation is to design a circuit  $U$  consisting of gates and oracles that approximates the time evolution up to an error  $\epsilon$  such that

$$\|U - e^{-iHt}\|_2 < \epsilon, \quad (1)$$

where  $\|\cdot\|_2$  is the spectral norm.




Figure 1: Hamiltonian simulation approximates the time evolution by a series of digital operations.

Quantum systems are fundamentally difficult to simulate; dynamics of quantum systems is a BQP-hard (or BQP complete for Hamiltonians with natural restrictions). Except for a few special cases, the complexity of the best known classical algorithms grows exponentially with the number of qubits. As such, simulation of quantum dynamics is a field where quantum computers can quickly outperform classical ones. In fact, the time evolution of quantum systems was the original application for quantum computers suggested by Feynman.

In the simplest scenario, we assume that the simulated Hamiltonians are given in the form  $H = \sum_{j=1}^m H_j$ , where each  $H_j$  is local and  $e^{-iH_j t}$  can be

implemented directly for arbitrary  $t$ . A common case is when these  $H_j$  are Pauli matrices.

The simplest quantum simulation algorithms rely on the Lie-Trotter formula

$$\lim_{r \rightarrow \infty} \left( e^{A/r} e^{B/r} \right)^r = \lim_{r \rightarrow \infty} \left( \left( 1 + \frac{A}{r} \right) \left( 1 + \frac{B}{r} \right) \right)^r \quad (2)$$

$$= \lim_{r \rightarrow \infty} \left( 1 + \frac{A+B}{r} + \frac{AB}{r^2} \right)^r \quad (3)$$

$$= \lim_{r \rightarrow \infty} \left( 1 + \frac{A+B}{r} \right)^r \quad (4)$$

$$= e^{A+B} \quad (5)$$

Since the individual terms in the Hamiltonian typically not commute, decomposing an exponential of a sum into a finite sum of exponentials will lead to errors.

**Exercise:** Prove that  $\left\| e^{A+B} - \left( e^{A/r} e^{B/r} \right)^r \right\| \in \mathcal{O} \left( \frac{\|A+B\| t^2}{r} \right)$ .

Next, we can recursively that for  $H = \sum_{j=1}^m H_j$ , one can decompose the evolution with respect to  $H$  into the evolution with respect to each  $H_j$  as

$$\tilde{U} = \left( e^{-iH_1 t/r} e^{-iH_2 t/r} \dots e^{-iH_m t/r} \right)^r + \mathcal{O}(\|H\| t^2/r). \quad (6)$$

Thus, if we are willing to tolerate error at most  $\epsilon$ , we need to perform  $\mathcal{O} \left( \frac{\|H\| t^2}{\epsilon} \right)$ .

Up to now, we assumed that we know how to implement each  $e^{-iH_j t}$  directly. Let us now show how to implement them in some simple cases.

In the simplest case,  $H_j$  is a Pauli matrix  $Z$  acting on the  $j$ th qubit. The Hamiltonian evolution is then a  $Z$ -rotation on the  $j$ th qubit.

$$e^{-itZ} = e^{-it} |0\rangle \langle 0| + e^{it} |1\rangle \langle 1| \quad (7)$$

Next, we show how to simulate a tensor product of  $Z$ s

$$H = Z_1 \otimes Z_2 \otimes \dots \otimes Z_n. \quad (8)$$

We first prove the following identity. For an arbitrary unitary  $U$ :

$$U e^{-iHt} U^\dagger = U \sum_{k=0}^{\infty} \frac{(-iHt)^k}{k!} U^\dagger \quad (9)$$

$$= UU^\dagger - iUHU^\dagger t + i^2 UH(UU^\dagger)Ht^2 U^\dagger - i^3 UH(UU^\dagger)H(UU^\dagger)HU^\dagger t^3 + \dots \quad (10)$$

$$= \sum_{k=0}^{\infty} \frac{(-iUHU^\dagger t)^k}{k!} \quad (11)$$

$$= e^{-iUHU^\dagger t}. \quad (12)$$




Figure 2: Simulating a tensor product of Pauli Zs.

We can implement  $e^{-iZ_1 \otimes \dots \otimes Z_n t}$  using the circuit in Fig. 2

The correctness of the circuit can be showed through induction. We already proved the first step in (7). In the inductive step, we conjugate an existing circuit by CNOTs. We can rewrite  $\text{CNOT} e^{-itZ_1 \otimes Z_2 \otimes \dots \otimes Z_{n-1}} \text{CNOT}$  as:

$$(\mathbb{I} \otimes |0\rangle\langle 0| + X \otimes |1\rangle\langle 1|) \left( \mathbb{I} \otimes e^{-itZ_1 \otimes \dots \otimes Z_{n-1}} \right) (\mathbb{I} \otimes |0\rangle\langle 0| + X \otimes |1\rangle\langle 1|) \quad (13)$$

$$= \otimes e^{-itZ_1 \otimes \dots \otimes Z_{n-1}} \otimes |0\rangle\langle 0| + X e^{-itZ_1 \otimes \dots \otimes Z_{n-1}} X \otimes |1\rangle\langle 1| \quad (14)$$

$$= e^{-itZ_1 \otimes \dots \otimes Z_{n-1}} \otimes |0\rangle\langle 0| + e^{itZ_1 \otimes \dots \otimes Z_{n-1}} \otimes |1\rangle\langle 1| \quad (15)$$

$$= e^{-itZ_1 \otimes Z_2 \otimes \dots \otimes Z_{n-1} \otimes Z_n} \quad (16)$$

Simulating other Paulis is possible by changing the basis. Recall that  $e^{-itUHU^\dagger} = Ue^{-itH}U^\dagger$ . We can then use the identities

$$X = \text{Had} Z \text{Had} \quad (17)$$

$$Y = S^\dagger \text{Had} Z \text{Had} S \quad (18)$$

We used Had instead of H to denote the Hadamard gate since we reserved the symbol H for Hamiltonians.

This allows us to simulate evolution according to any Pauli. For example, we can simulate the evolution according to the  $H = X \otimes Y \otimes Z$  according to the circuit in Fig. 3.




Figure 3: X and Y Paulis can be simulated by a change of basis. The circuit above depicts the simulation according to  $e^{-itX \otimes Y \otimes Z}$ .

Now we know how to simulate any Hamiltonian that is a sum of Paulis. The complexity of the algorithm for simulating a Hamiltonian  $H = \sum_{l=0}^{L-1} P_l$  where

$P$  is a Pauli on  $n$  qubits is

$$\mathcal{O}\left(\frac{Lt^2n}{\epsilon}\right) \quad (19)$$

However, there is still a need for more efficient algorithms. First, the number of terms  $L$  needed to decompose a Hamiltonian into a sum of Paulis can be exponentially large. Second, the scaling in terms of  $t$  and  $\epsilon$  is quite poor.

Other Hamiltonian simulation algorithms allow for different decomposition of Hamiltonians. A popular one is decomposing a sparse matrix into 1-sparse matrices which can be then simulated directly. Another one is decomposition into a linear combination of unitaries (LCU) which a generalization of the Pauli decompositions.

**Exercise:** Let  $\rho, \sigma$  be density matrices,  $S$  the SWAP operator and  $\text{Tr}_p$  partial trace over the first variable. Show that

$$\text{Tr}_p[e^{-iS\Delta}\rho \otimes \sigma e^{iS\Delta}] = e^{-i\rho\Delta}\sigma e^{i\rho\Delta} + \mathcal{O}(\Delta^2)$$

In these more general cases the Hamiltonian is accessed through an oracle. One type of oracular access is particularly common when the Hamiltonian (in a computational basis) is given by a sparse matrix. We say a Hamiltonian is *row-d-sparse* if each row has at most  $d$  non-zero entries. If there is an efficient procedure to locate these entries we moreover say that the Hamiltonian is *row-computable*. In this case, one can efficiently construct oracles

$$O_{\text{loc}} |r, k\rangle = |r, k \oplus l\rangle \quad (20)$$

$$O_{\text{val}} |r, l, z\rangle = |r, l, z \oplus H_{r,l}\rangle. \quad (21)$$

Oracle  $O_{\text{loc}}$  locates the position  $l$  of the  $k$ -th non-zero element in row  $r$ . The oracle  $O_{\text{val}}$  then gives the value of the matrix element  $H_{r,l}$ . We compute the cost of algorithms in terms of the number of queries to these oracles.

It is possible to construct different oracles. Any Hermitian matrix can be decomposed into a sum of unitaries

$$H = \sum_{l=0}^{L-1} \alpha_l V_l, \quad (22)$$

where for each  $l$ ,  $\alpha_l \leq 0$  and  $H_l$  is a unitary matrix  $\|H_l\| = 1$ . This decomposition can be efficiently implemented for sparse Hamiltonians. The coefficients  $\alpha_l$  and unitaries  $H_l$  can be accessed through oracles

$$O_\alpha |l, z\rangle = |l, z \oplus \alpha_l\rangle \quad (23)$$

$$O_{V_l} |l, \psi\rangle = V_l |l, \psi\rangle, \quad (24)$$

or, in some cases, described classically.

## 1.1 Further reading

Efficient quantum algorithms for simulating sparse Hamiltonians

G. Ahokas, R. Cleve, B.C. Sanders, D.W. Berry

arXiv:quant-ph/0508139

Simulating Hamiltonian dynamics with a truncated Taylor series

D. W. Berry, A. M. Childs, Richard Cleve, R. Kothari, R. D. Somma

arXiv:1412.4687

Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics

A Gilyén, Y Su, GH Low, N Wiebe

arXiv:1806.01838

## 2 Linear Systems

An important application of Hamiltonian simulation is “solving” systems of linear equations. The definition of solving is different for quantum and classical algorithms. While the classical algorithm outputs the solution of the linear system, quantum algorithms encode the solution into a quantum state. We will first review the classical approaches to solving systems of equations and then explain the corresponding quantum algorithms.

A system of equations can be written in a standardized form

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ &\vdots \\ a_{n1}x_1 + a_{n2}x_2 + \cdots + a_{nn}x_n &= b_n, \end{aligned} \tag{25}$$

or a matrix equation

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{26}$$

where  $\mathbf{A}$  encode the coefficients in the equations, the vector  $\mathbf{x}$  are the unknowns and  $\mathbf{b}$  stores the right hand side of equations. The solution then can be written as

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b} \tag{27}$$

if the matrix  $\mathbf{A}$  is indeed invertable. We will for now assume that this is the case and consider the linear system problem to be equivalent to matrix inversion.

One of the most well-known algorithms is Gaussian elimination with complexity  $\mathcal{O}(N^3)$ . While the complexity is polynomial in the number of equations and therefore “efficient”, cubic runtime is impractical for large systems. Therefore, there is a lot of interest in finding more efficient exact and approximation algorithms.

The first quantum linear systems algorithm is referred to as HHL after the initials of its creators - Harrow, Hassidim and Lloyd. The linear system

algorithm makes efficient use of phase estimation. It is sufficient to consider Hermitian matrix; if the matrix  $A$  is not hermitian, a Hermitian matrix can be created as  $H = \begin{pmatrix} \cdot & A \\ A^\dagger & \cdot \end{pmatrix}$  and take the right hand side into  $\begin{pmatrix} b \\ 0 \end{pmatrix}$ . The solution will this new system will then be  $\begin{pmatrix} 0 \\ x \end{pmatrix}$ . Without the lost of generality, we will from now on assume that  $A$  is hermitian.

Let  $|\alpha_i\rangle$  and  $\lambda_i$  be the eigenvectors and eigenvalues of  $A$  and  $|b\rangle = \sum_i \beta_i |\alpha_i\rangle$  be a decomposition of the right-hand side into this basis.

Using Hamiltonian simulation (with  $t=1$ ) we and phase estimation we can prepare

$$|b\rangle |0 \dots 0\rangle |0\rangle \rightarrow \sum_i \beta_i |\alpha_i\rangle |\tilde{\lambda}_i\rangle |0\rangle. \quad (28)$$

Next, we apply a rotation conditioned on  $|\tilde{\lambda}_i\rangle$  to prepare

$$\sum_i \beta_i |\alpha_i\rangle |\tilde{\lambda}_i\rangle \left( \sqrt{1 - \frac{C^2}{\tilde{\lambda}_i^2}} |0\rangle + \frac{C}{\tilde{\lambda}_i} |1\rangle \right), \quad (29)$$

where  $C = \mathcal{O}(1/\kappa)$ . The phase estimation is the uncomputed, erasing the second register, and we are left with

$$\left( \sqrt{1 - \frac{C^2}{\tilde{\lambda}_i^2}} \sum_i \beta_i |\alpha_i\rangle |0 \dots 0\rangle |0\rangle + \frac{C}{\tilde{\lambda}_i} \sum_i \beta_i |\alpha_i\rangle |0 \dots 0\rangle |1\rangle \right). \quad (30)$$

In case when the last qubit is in state 1, the resulting state on the first register is proportional to  $\sum_i \lambda_i^{-1} \beta_i |\lambda_i\rangle$

Overall, the complexity of the algorithm is  $\tilde{\mathcal{O}}\left(\log n \frac{s^2 \kappa^2}{\epsilon}\right)$ .

While HHL achieves exponential improvement in  $n$  when naively compared to classical algorithms, its scaling in the error  $\epsilon$  and the conditioning number  $\kappa$  is not optimal.

However, this comparison is not completely fair. First, the algorithm is limited to matrices that can be effectively encoded in a quantum computer which are usually either sparse or low-rank matrices. Second, the state  $b$  also needs to be prepared effectively on a quantum computer. Lastly, the output of HHL is a quantum state and outputting it (i.e., writing it on a paper) would result in an exponential overhead and worse performance than most classical algorithms. Instead, the task of the quantum computer is often only to sample from the solution. A recent result managed to "de-quantize" many applications of HLL and showed that an algorithm for "solving" linear systems with a low-rank matrix and sampling from the solution can be simulated classically.

## 2.1 Further reading

Quantum algorithm for solving linear systems of equations

AW Harrow, A Hassidim, S Lloyd  
arXiv:0811.3171

Quantum Machine Learning Algorithms: Read the Fine Print S Aaronson  
<https://www.scottaaronson.com/papers/qml.pdf>

Sampling-based sublinear low-rank matrix arithmetic framework for de-quantizing quantum machine learning

N Chia, A Gilyén, T Li, H Lin, E Tang, C Wang  
arXiv:1910.06151

### 3 Quantum Machine Learning (QML)

HLL can be applied to various quantum machine learning problems. QML build on the intuition that most of machine learning is some form of linear algebra in high-dimensional spaces and quantum computing is good at performing algebra at high-dimensional spaces<sup>1</sup>. However, we will see that that obtaining quantum speedup for machine learning algorithms is not so straightforward.

#### 3.1 Least-squares fitting

The first application is to apply HLL algorithm for data fitting. Our goal is to find a parameters  $\lambda$  such that a parametric function  $f(\mathbf{x}, \lambda)$  will provide an optimal fit for  $N$  datapoints  $\{x_i, y_i\}$ .

The function is constrained to be linear in the fit parameters  $\lambda \in \mathbb{C}^M$ , but it can be non-linear in  $\mathbf{x}$ . For example, we can encode  $f(\mathbf{x}, \lambda) = a + b\mathbf{x} + c\mathbf{x}^2$  as  $\lambda = (a, b, c)$  and  $\mathbf{x} = (1, \mathbf{x}, \mathbf{x}^2)$ .

To goal of the algorithm is then find the parameters  $\lambda$  to minimize

$$E = \sum_{i=1}^N |f(x_i, \lambda) - y_i|^2 = \|\mathbf{F}\lambda - \mathbf{y}\|^2. \quad (31)$$

Here we defined  $F_{ij} = f_j(x_i)$  to be an  $N \times M$  matrix. Note that  $\mathbf{F}$  is generally not invertable, however, we can assume that  $\frac{1}{\kappa} \leq \|\mathbf{F}^\dagger \mathbf{F}\| < 1$ . One can then show that the minimum of (31) is achieved by

$$\lambda = \mathbf{F}^+ \mathbf{y} = \left(\mathbf{F}^\dagger \mathbf{F}\right)^{-1} \mathbf{F} \mathbf{y}, \quad (32)$$

where  $\mathbf{F}^+$  is Moore–Penrose pseudoinverse.

The quantum version of least squares fitting assumes that  $\mathbf{y}$  is encoded into a quantum states  $|\mathbf{y}\rangle = \sum_{j=M+1}^{M+N} y_j |j\rangle / \text{norm} \mathbf{y}$ . Similarly, the output  $|\lambda\rangle$  will be also given only as a quantum state. Additionally, we have an oracular access to

---

<sup>1</sup>From Seth Lloyd, paraphrased




Figure 4: SVM for 2-dimensional data. The red line maximizes the separation between the classes. Source: Wikipedia

$\tilde{F} = \begin{pmatrix} \cdot & F \\ F^\dagger & \cdot \end{pmatrix}$ . For an efficient algorithm,  $\tilde{F}$  must belong into a class of simulable Hamiltonians such as sparse or low-rank ones.

The algorithm for least square fitting is then a technical modification of the HLL algorithm. First, we need to prepare  $F|y\rangle$ . This can be accomplished using HHL but though Hamiltonian evolution and phase estimation just as in HLL but instead of diving by the eigenvalues we will multiply by them. The inverse of  $F^\dagger F$  is accomplished though HHL as well. The resulting query complexity of the algorithm is  $\tilde{O}(\log(N)s^4\kappa^6/\epsilon)$  where  $s$  is the sparsity and the notation  $\tilde{O}$  suppresses log-factors.

### 3.2 Quantum Support Vector Machine

Support vector machine (SVM) is a supervised machine learning algorithm for data classification. The algorithm is given a set of data where each data point belongs to one of two categories (green or blue). The goal of the algorithm is to find a hyperplane that would separate the green points from the blue ones. Since there are many such hyperplanes, we aim to find one that separates them with the largest margin, see Fig. 4. This is an example of a non-probabilistic binary linear classifier and although SVMs can be generalized to find non-linear classifiers, we will focus on this simple example.

We know from analytic geometry that a hyperplane is a set of points  $\mathbf{x}$  that satisfy

$$\mathbf{w} \cdot \mathbf{x} - b = 0, \quad (33)$$

where  $\mathbf{w}$  is the normal vector. Note that we do not require  $w$  to be normalized –



its norm will determine the side of the boundary. We will require that all the blue points  $\mathbf{x}$  satisfy

$$\mathbf{w} \cdot \mathbf{x} - b > 1, \tag{34}$$

and for each the green point

$$\mathbf{w} \cdot \mathbf{x} - b < -1, \tag{35}$$

as in Fig. ?? . If we use  $y = +1$  to label all the blue data and  $y = -1$ , we can consolidate (34) and (35) as

$$y_i(\mathbf{w} \cdot \mathbf{x}_i - b) > 1, \quad \forall (x_i, y_i) \in \text{training set}, . \tag{36}$$

The goal of the training is then to minimize  $\|\mathbf{w}\|$  subject to the constrain (36).

To turn SVM into a quantum algorithm, we use the usual tricks, turning vectors into quantum states, and reformulate (36) and reformulating SVM as a linear system that can be solved using HLL.

### 3.3 The weaknesses of quantum machine learning

We saw two algorithms that appear to give exponential speedup over their classical counterparts. Indeed, in our setting, the runtimes of the algorithms are faster than even writing the complete input. This is because the assumptions about the input and output of the quantum algorithms are much stronger than what is considered in the classical case. Let us reiterate the assumptions: first, we assume that the input data is encoded as amplitudes of a quantum state. While we have such an encoding is possible, preparing such state, in general, requires exponentially many gates. Second, we assumed that any matrices can be embedded into easy to simulate Hamiltonian and easily accessible. This means that the exponential speedup exists ONLY with respect to an oracle that encodes such a Hamiltonian. Lastly, we are only required to output the solution as a quantum state instead of fully writing it out.

The input problem is abstracted by a quantum random memory (QRAM). QRAM can return data as

$$\text{QRAM } |j\rangle |0\dots 0\rangle = |j\rangle |\psi_j\rangle \tag{37}$$

and can operate in superposition. See <https://youtu.be/IicCWK2D7sg> for QRAM construction. If the Hermitian matrix encoding the data is low-rang, QRAM allows for efficient implementation of QML algorithms

Surprisingly, Ewin Tang (followed by others) managed to dequantize multiple quantum algorithm that relied on QRAM constructions. Her approach used sampling and randomized algebra techniques. As such, there are currently no known QML algorithms with provable, non-oracular exponential speedup.

### 3.4 Further Reading

Quantum machine learning

J Biamonte, P Wittek, N Pancotti, P Rebentrost, N Wiebe, S Lloyd  
arXiv:1611.09347

A quantum-inspired classical algorithm for recommendation systems

E Tang  
arXiv:1807.04271v3

Key questions for the quantum machine learner to ask themselves

N Wiebe  
<https://iopscience.iop.org/article/10.1088/1367-2630/abac39/meta>