CS 4134 – Quantum Computation and Information Processing

Sumeet Khatri

January 20, 2026

Course format

- ▶ Lectures by me. Video recordings will be posted to Canvas after class.

- ▶ Assignments for you to complete (on a weekly basis).

- ▶ Five in-class written exams throughout the semester — top four will be counted.

- ▶ Prerequisites: background in linear algebra and probability/statistics is strongly recommended. But we will cover the basics!

- ▶ We will use Piazza for questions and discussions about assignments, etc.

- ▶ We have four TAs for the class — they will hold office hours and be available on Piazza.

Exam and date	
Exam 1	February 12, 2026
Exam 2	March 5, 2026
Exam 3	March 26, 2026
Exam 4	April 16, 2026
Exam 5	May 7, 2026

Final grade	Scheme 1	Scheme 2
Exams	80%	100%
Assignments	20%	0%

★ Your final grade will be the highest among the two schemes.

Topics outline

- Part I: Mathematics and phenomena of quantum (information) theory

- Part II: Quantum circuits and algorithms

- Part III: Noise in quantum computing, error correction

Recommended texts

⋆ *Quantum Computation and Quantum Information*, by Michael Nielsen and Isaac Chuang.

⋆ *An Introduction to Quantum Computing, by Phillip Kaye*, Raymond Laflamme and Michele Mosca.

⋆⋆ *Introduction to Classical and Quantum Computing*, by Thomas G. Wong. **(Nice and gentle.)**

▶ *Quantum Information: From Foundations to Quantum Technology Applications*, by Dagmar Bruss and Gerd Leuchs (editors). **(Very broad in scope, discusses applications.)**

▶ *The Theory of Quantum Information*, by John Watrous.

▶ *Classical and Quantum Computation*, by A. Yu. Kitaev, A. H. Shen, M. N. Vyalyi. **(Quite advanced!)**

▶ *Principles of Quantum Communication Theory: A Modern Approach*, by Sumeet Khatri and Mark M. Wilde. Available at `https://arxiv.org/abs/2011.04672`.

Quantum mechanics

- In short, the study of matter on small scales. (Atoms/molecules, particles of light (photons).)

Quantum mechanics

▶ In short, the study of matter on small scales. (Atoms/molecules, particles of light (photons).)
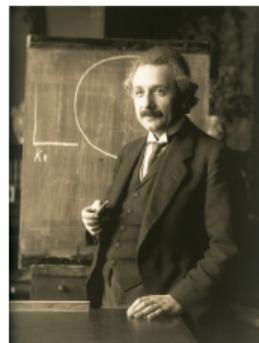
Foundational discoveries

Planck
Blackbody radiation
$E = h\nu$
(Discrete energies)

Heisenberg
$\Delta X \Delta P \geq \frac{\hbar}{2}$
(Uncertainty relation)

Schrödinger
$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H |\psi(t)\rangle$
(Dynamics)

Einstein
Photoelectric effect
EPR paradox
(Entanglement)

Born
Born rule
(Measurements)

Quantum mechanics

► In short, the study of matter on small scales. (Atoms/molecules, particles of light (photons).)
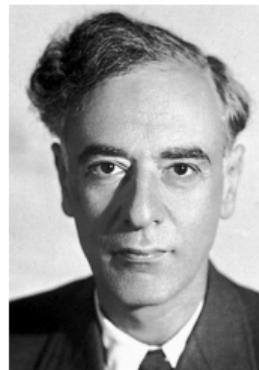
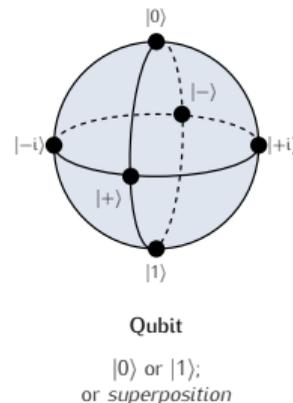### Axiomatic mathematical formulation
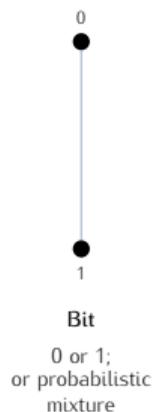

Dirac   von Neumann   Landau

We will go through this in detail later!

Quantum mechanics → quantum information

▶ Use quantum properties of atoms/molecules and photons as *carriers of information*.

▶ For example, energy levels of atoms/molecules, polarizations of light/photons.

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle,$$
$$\alpha, \beta \in \mathbb{C},$$
$$|\alpha|^2 + |\beta|^2 = 1$$

$$|\pm\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$$
$$|\pm i\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle \pm i|1\rangle)$$

Bit
0 or 1;
or probabilistic
mixture

Qubit
$|0\rangle$ or $|1\rangle$;
or *superposition*

▶ To extract information from a qubit, we have to *measure* it.

▶ The measurement outcomes are not deterministic! ("0" and "1" occur with certain probabilities) — i.e., every time you measure, you can generally get a different outcome!

▶ Qubits can be *entangled* — this allows for teleportation!

★ To understand these things, we need need a solid grasp of linear algebra and basic probability theory and statistics!

No-cloning theorem

The Concept of Transition in Quantum Mechanics

James L. Park

Department of Physics, Washington State University, Pullman, Washington

A single quantum cannot be cloned

W. K. Wootters*

Center for Theoretical Physics, The University of Texas at Austin, Austin, Texas 78712, USA

W. H. Zurek

Theoretical Astrophysics 130–33, California Institute of Technology, Pasadena, California 91125, USA

COMMUNICATION BY EPR DEVICES

D. DIEKS
Fysisch Laboratorium, Rijksuniversiteit Utrecht, Utrecht, The Netherlands

A recent proposal to achieve faster-than-light communication by means of an EPR-type experimental set-up is examined. We demonstrate that such superluminal communication is not possible. The crucial role of the linearity of the quantum mechanical evolution laws in preventing causal anomalies is stressed.

There does not exist a unitary (linear) operation that can copy an arbitrary quantum state.

Bell's theorem

ON THE EINSTEIN PODOLSKY ROSEN PARADOX*

J. S. BELL[†]

*Department of Physics, University of Wisconsin, Madison, Wisconsin*

(Received 4 November 1964)

Correlations in quantum mechanics do not generally correspond to a local-hidden-variable model.

Nobel Prize 2022!

Quantum information science and quantum technologies

► Harness superposition and entanglement to do certain tasks better/faster. (e.g., factoring, simulation)
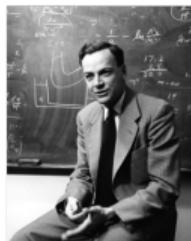► Also discover *new* things. (e.g., teleportation, QKD)
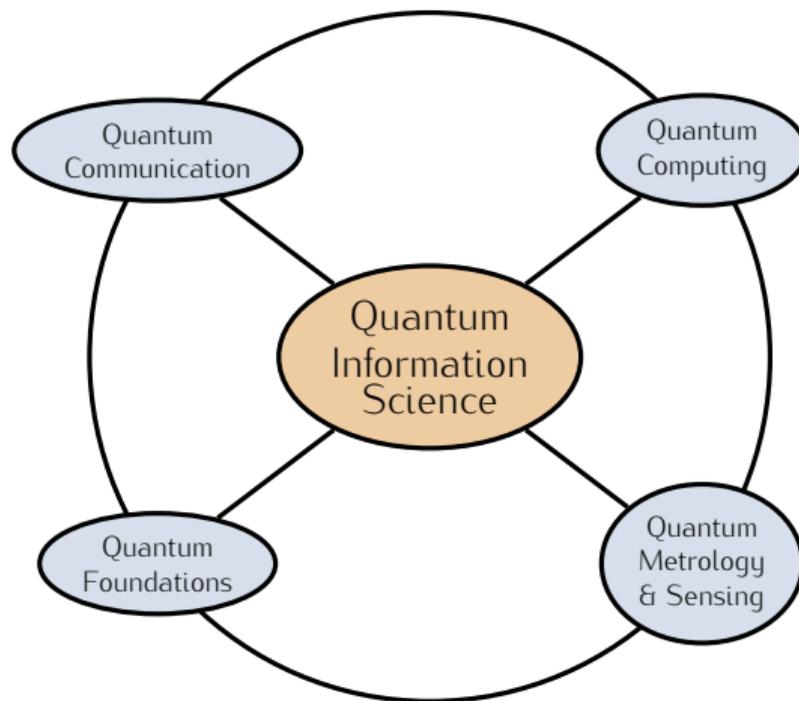
International Journal of Theoretical Physics, Vol. 21, Nos. 6/7, 1982

Simulating Physics with Computers

Richard P. Feynman

Department of Physics, California Institute of Technology, Pasadena, California 91107

Received May 7, 1981

Proc. R. Soc. Lond. A **400**, 97–117 (1985)
Printed in Great Britain

Quantum theory, the Church–Turing principle and the universal quantum computer

BY D. DEUTSCH

Department of Astrophysics, South Parks Road, Oxford OX1 3RQ, U.K.

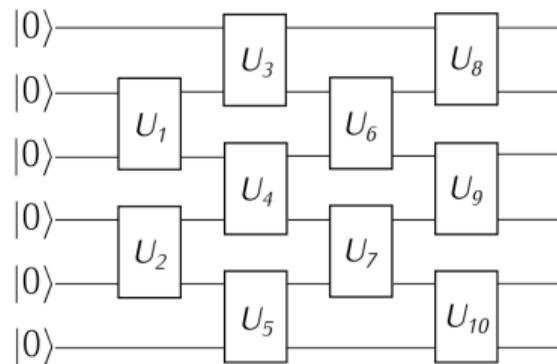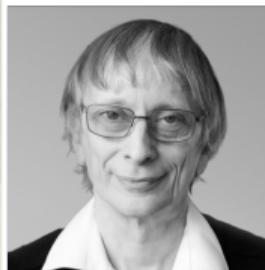(Communicated by R. Penrose, F.R.S. – Received 13 July 1984)

Proc. R. Soc. Lond. A **425**, 73–90 (1989)
Printed in Great Britain

Quantum computational networks

BY D. DEUTSCH

Oxford University Mathematical Institute, 24–29 St Giles, Oxford OX1 3LB, U.K.

(Communicated by R. Penrose, F.R.S. – Received 8 July 1988)

★ Potential for speedups in simulating large quantum systems.

★ Known speedups for factoring (Shor) and database search (Grover).

Algorithms for Quantum Computation:
Discrete Logarithms and Factoring

Peter W. Shor   (1994)
AT&T Bell Labs
Room 2D-149
600 Mountain Ave.
Murray Hill, NJ 07974, USA

**POLYNOMIAL-TIME ALGORITHMS FOR PRIME FACTORIZATION
AND DISCRETE LOGARITHMS ON A QUANTUM COMPUTER[*]**

PETER W. SHOR[†]

Prime factorization in polynomial time!

A fast quantum mechanical algorithm for database search
Lov K. Grover
3C-404A, AT&T Bell Labs
600 Mountain Avenue   (1996)
Murray Hill NJ 07974
lkg@mhcnet.att.com

Quadratic speedup for search problems!
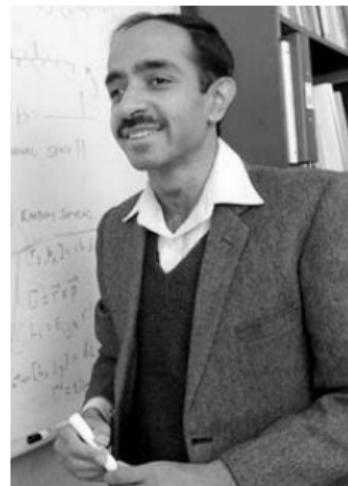
Quantum information science and quantum technologies

▶ Harness superposition and entanglement to do certain tasks better/faster. (e.g., factoring, simulation)
▶ Also discover *new* things. (e.g., teleportation, QKD)
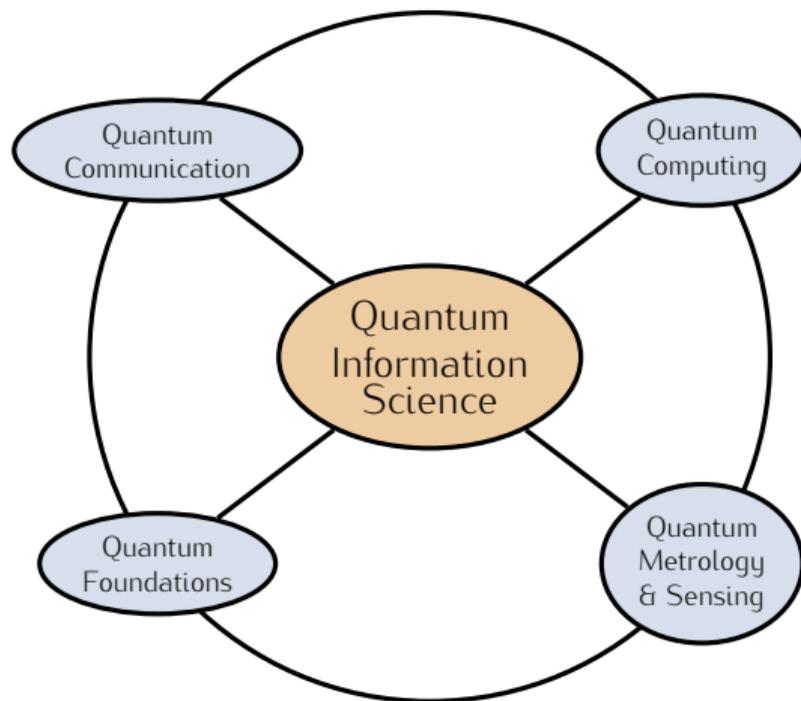
★ Using quantum systems and quantum strategies to send bits and qubits.

Alice

Bob

Ideal quantum channel from Alice to Bob.

⋆ Using quantum systems and quantum strategies to send bits and qubits.

Alice

Bob
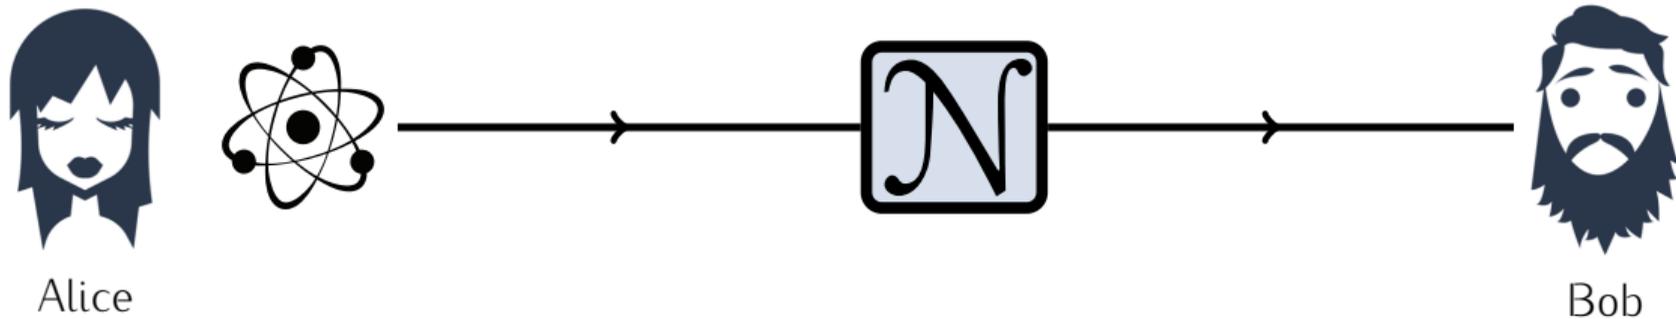
Noisy quantum channel from Alice to Bob, models imperfections in the transmission medium.

Teleportation: shared entanglement + classical communication = transmission of an arbitrary quantum state

PHYSICAL REVIEW
LETTERS

VOLUME 70	29 MARCH 1993	NUMBER 13

Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels

Charles H. Bennett,[1] Gilles Brassard,[2] Claude Crépeau,[2],[3]
Richard Jozsa,[2] Asher Peres,[4] and William K. Wootters[5]

No physical transmission of quantum systems! Only classical communication!

Teleportation: shared entanglement + classical communication = transmission of an arbitrary quantum state
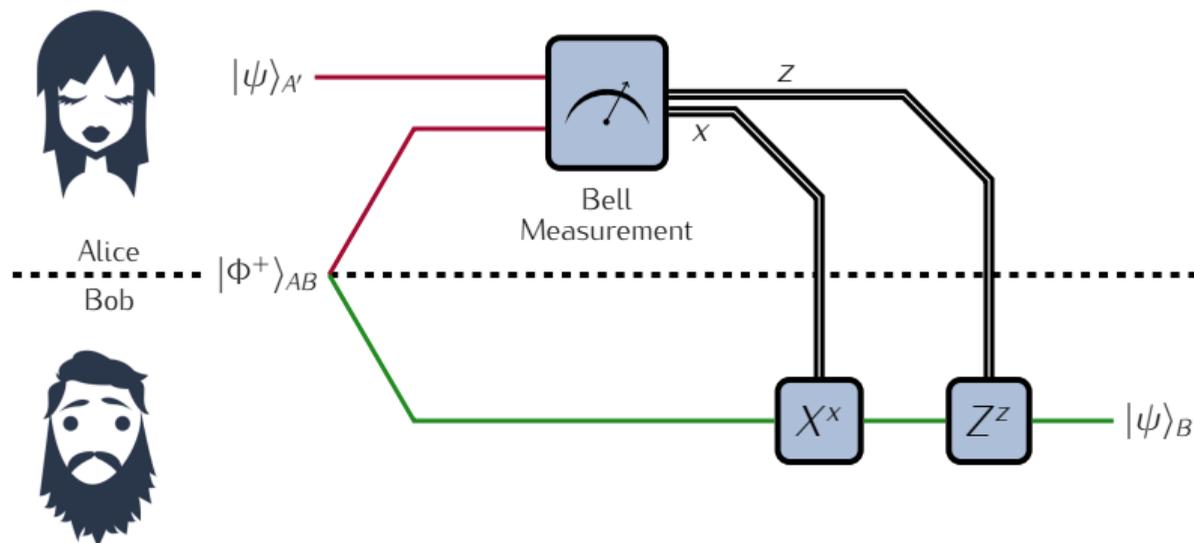
No physical transmission of quantum systems! Only classical communication!

Quantum key distribution: private classical communication with quantum strategies/resources

QUANTUM CRYPTOGRAPHY: PUBLIC KEY DISTRIBUTION AND COIN TOSSING

Charles H. Bennett (IBM Research, Yorktown Heights NY 10598 USA)
Gilles Brassard (dept. IRO, Univ. de Montreal, H3C 3J7 Canada)

International Conference on Computers, Systems & Signal Processing   Bangalore, India   December 10-12, 1984

Quantum key distribution: private classical communication with quantum strategies/resources
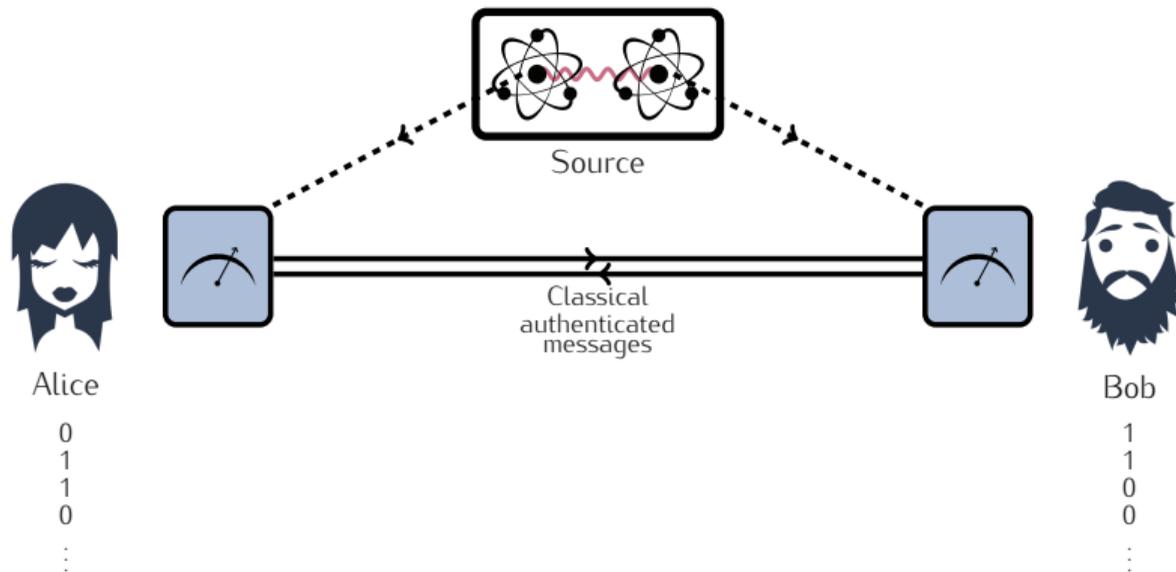
More recent developments

★ Making actual qubits is *hard*.

★ Theoretical protocols do not perform as expected in practice – there are errors!
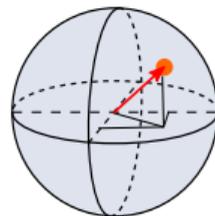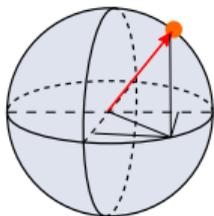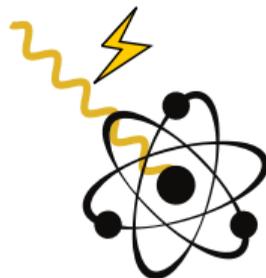
* Causes of noise: interaction with the environment, imperfections in the medium.
* Mathematical model: completely positive trace–preserving map (comes from the Schrödinger equation).
* Noise means that theoretical schemes do not necessarily work exactly as described in practice!

Consequences of noise for quantum computation

$\star$ So far, we assumed that every gate U_i is unitary.

$\star$ In practice, there are two types of errors:

  $\blacktriangleright$ Coherent errors: $\widetilde{U}_i$ instead of U_i.

  $\blacktriangleright$ Incoherent errors: A *quantum channel* $\mathcal{N}_i$ instead of U_i
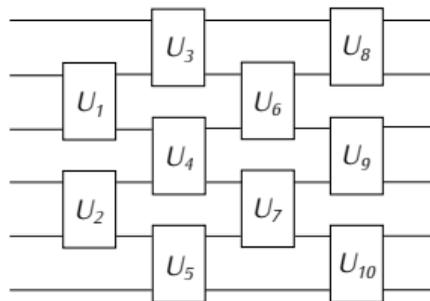
$\star$ **How much noise can we tolerate while still doing a useful computation?**

Consequences of noise for quantum computation

★ So far, we assumed that every gate U_i is unitary.

★ In practice, there are two types of errors:

  ▶ Coherent errors: $\widetilde{U}_i$ instead of U_i.

  ▶ Incoherent errors: A *quantum channel* $\mathcal{N}_i$ instead of U_i
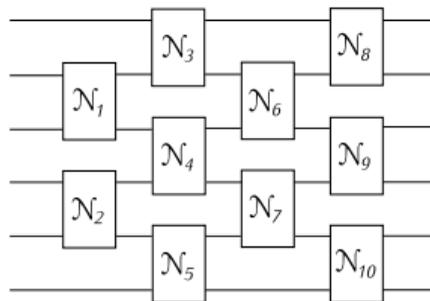
★ **How much noise can we tolerate while still doing a useful computation?**

PHYSICAL REVIEW A

ATOMIC, MOLECULAR, AND OPTICAL PHYSICS

RAPID COMMUNICATIONS

Scheme for reducing decoherence in quantum computer memory

Peter W. Shor*
AT&T Bell Laboratories, Room 2D-149, 600 Mountain Avenue, Murray Hill, New Jersey 07974
(Received 17 May 1995)

Fault-Tolerant Quantum Computation

Peter W. Shor
AT&T Research
Room 2D-149
600 Mountain Ave.
Murray Hill, NJ 07974, USA
shor@research.att.com

Fighting noise and errors

▶ Analogous to (classical) error correction,
$0 \mapsto 000$, $1 \mapsto 111$.

$$|0\rangle \mapsto |0, 0, 0\rangle, \quad |1\rangle \mapsto |1, 1, 1\rangle.$$

▶ Fault-tolerant schemes allow for error-free
computation, in principle.

▶ But this requires many good-quality qubits:
~ 1000 physical qubits per logical qubit.

From [https://www.osti.gov/servlets/purl/1593593]

Fighting noise and errors

- Analogous to (classical) error correction,
  $0 \mapsto 000$, $1 \mapsto 111$.

  $$|0\rangle \mapsto |0, 0, 0\rangle, \quad |1\rangle \mapsto |1, 1, 1\rangle.$$

- Fault–tolerant schemes allow for error–free
  computation, in principle.

- But this requires many good-quality qubits:
  ~ 1000 physical qubits per logical qubit.

From [https://www.osti.gov/servlets/purl/1640593]

FIG. 1 (color online). A surface code logical qubit. Stabilizers
$ZZZZ$ ($XXXX$) are associated with the data qubits (open circles)
around each face (vertex). Syndrome qubits (dots) measure
stabilizers using the indicated sequences of gates. Logical op-
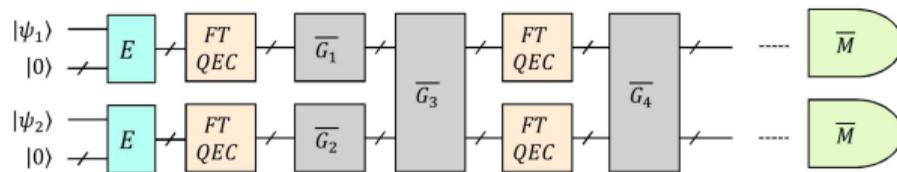erators Z_L, X_L connect opposing boundaries.

From [PRA **86**, 032324 (2012)]

Fighting noise and errors

- Analogous to (classical) error correction, $0 \mapsto 000$, $1 \mapsto 111$.

$$|0\rangle \mapsto |0, 0, 0\rangle, \ |1\rangle \mapsto |1, 1, 1\rangle.$$

- Fault–tolerant schemes allow for error–free computation, in principle.

- But this requires many good-quality qubits: ~ 1000 physical qubits per logical qubit.

- We currently have:
  - Limited number of qubits (≈ 100, approaching 1000).
  - Limited connectivity between qubits.
  - Limited circuit depth due to noise.

From [https://www.osti.gov/servlets/purl/1640593]

FIG. 1 (color online). A surface code logical qubit. Stabilizers $ZZZZ$ ($XXXX$) are associated with the data qubits (open circles) around each face (vertex). Syndrome qubits (dots) measure stabilizers using the indicated sequences of gates. Logical operators Z_L, X_L connect opposing boundaries.

From [PRA **86**, 032324 (2012)]

▶ Analogous to (classical) error correction, $0 \mapsto 000$, $1 \mapsto 111$.

$$|0\rangle \mapsto |0, 0, 0\rangle, \ |1\rangle \mapsto |1, 1, 1\rangle.$$

▶ Fault–tolerant schemes allow for error–free computation, in principle.

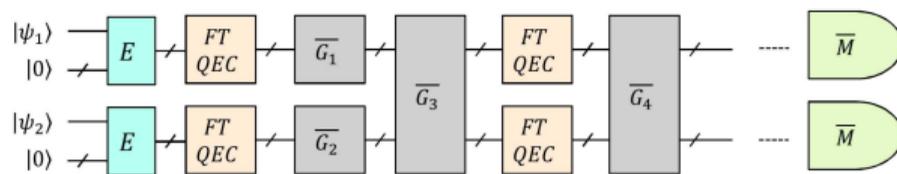▶ But this requires many good-quality qubits: ~ 1000 physical qubits per logical qubit.

▶ We currently have:

　▶ Limited number of qubits (≈ 100, approaching 1000).

　▶ Limited connectivity between qubits.

　▶ Limited circuit depth due to noise.

★ Much current work is devoted to discovering and implementing error-correcting codes!
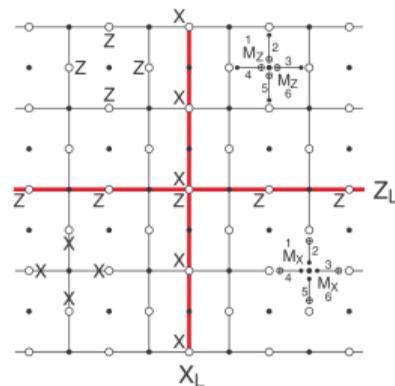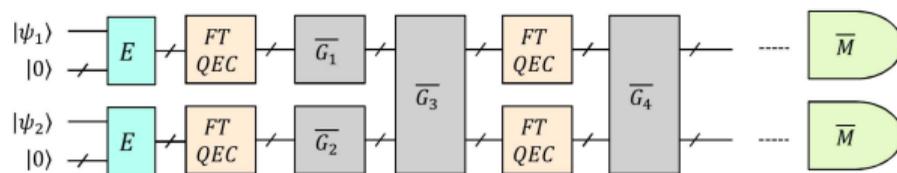

From [https://www.osti.gov/servlets/purl/1640593]
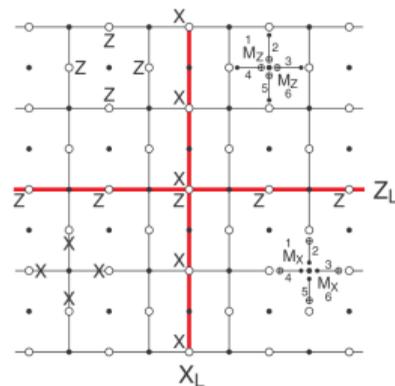
FIG. 1 (color online). A surface code logical qubit. Stabilizers $ZZZZ$ ($XXXX$) are associated with the data qubits (open circles) around each face (vertex). Syndrome qubits (dots) measure stabilizers using the indicated sequences of gates. Logical operators Z_L, X_L connect opposing boundaries.

From [PRA **86**, 032324 (2012)]

Status of current quantum computers

[https://www.ibm.com/quantum/technology]
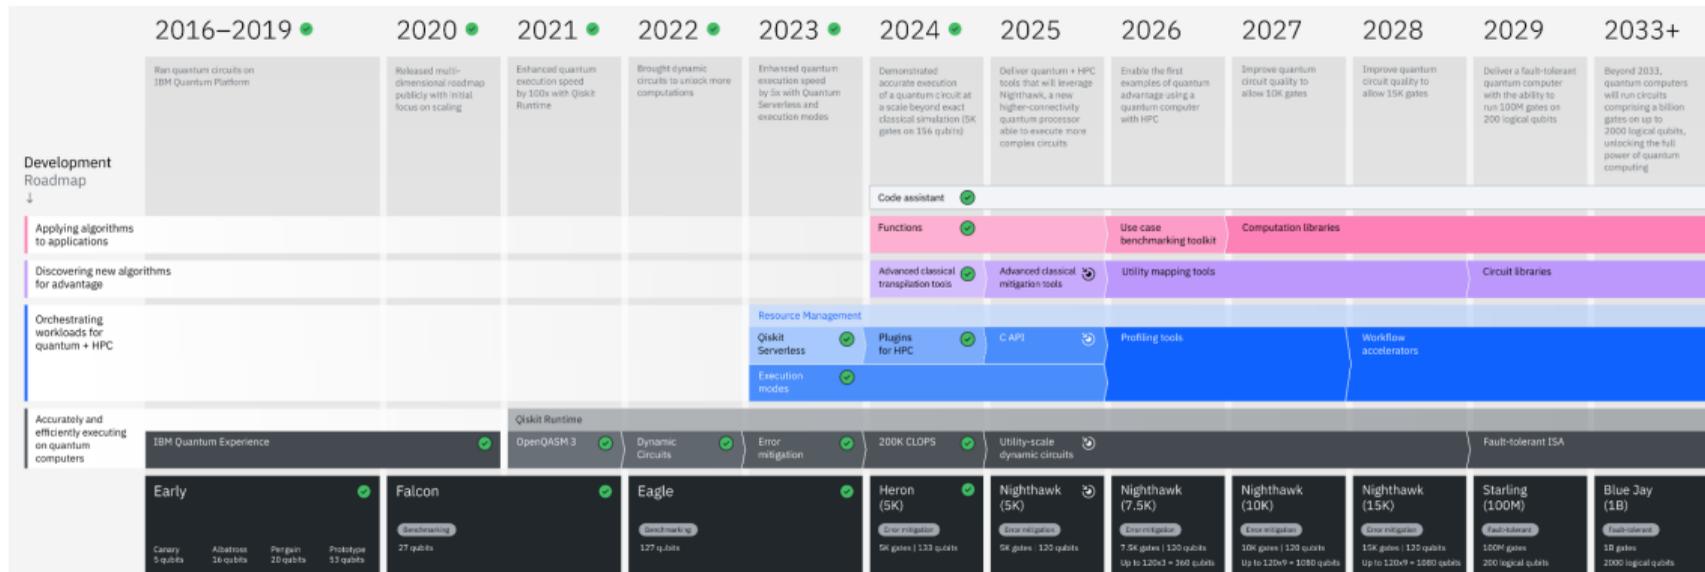
Quantum Computing in the NISQ era and beyond

John Preskill

Institute for Quantum Information and Matter and Walter Burke Institute for Theoretical Physics,
California Institute of Technology, Pasadena CA 91125, USA
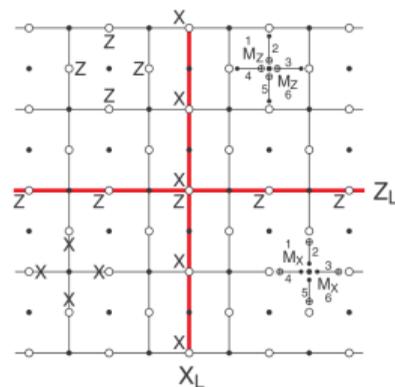30 July 2018

Noisy Intermediate-Scale Quantum (NISQ) technology will be available in the near future. Quantum computers with 50-100 qubits may be able to perform tasks which surpass the capabilities of today's classical digital computers, but noise in quantum gates will limit the size of quantum circuits that can be executed reliably. NISQ devices will be useful tools for exploring many-body quantum physics, and may have other useful applications, but the 100-qubit quantum computer will not change the world right away — we should regard it as a significant step toward the more powerful quantum technologies of the future. Quantum technologists should continue to strive for more accurate quantum gates and, eventually, fully fault-tolerant quantum computing.

[arXiv:1801.00862]

It's a remarkable claim — one of the most amazing ideas I've encountered in my scientific life — that there is a distinction between problems that are classically hard and problems that are quantumly hard. And it is a compelling challenge to understand better what problems are classically hard but quantumly easy [4, 5]. We should recognize in particular that the power of a quantum computer is not unlimited. We don't expect, for example, that a quantum computer will be able to solve efficiently the hard instances of NP-hard problems like the traveling salesman problem. For such hard combinatorial search problems we probably can't do much better than exhaustively searching for a solution. Quantum computers can speed up exhaustive search [6], but only modestly [7], so NP-hard problems are likely to be quantumly hard as well as classically hard.

[arXiv:1801.00862]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

QUANTUM COMPUTING AND
THE ENTANGLEMENT FRONTIER

JOHN PRESKILL

Institute for Quantum Information and Matter
California Institute of Technology
Pasadena, CA 91125, USA

Quantum information science explores the frontier of highly complex quantum states, the "entanglement frontier." This study is motivated by the observation (widely believed but unproven) that classical systems cannot simulate highly entangled quantum systems efficiently, and we hope to hasten the day when well controlled quantum systems can perform tasks surpassing what can be done in the classical world. One way to achieve such "quantum supremacy" would be to run an algorithm on a quantum computer which solves a problem with a super-polynomial speedup relative to classical computers, but there may be other ways that can be achieved sooner, such as simulating exotic quantum states of strongly correlated matter. To operate a large scale quantum computer reliably we will need to overcome the debilitating effects of decoherence, which might be done using "standard" quantum hardware protected by quantum error-correcting codes, or by exploiting the nonabelian quantum statistics of anyons realized in solid state systems, or by combining both methods. Only by challenging the entanglement frontier will we learn whether Nature provides extravagant resources far beyond what the classical world would allow.

Rapporteur talk at the 25th Solvay Conference on Physics
"The Theory of the Quantum World"
Brussels, 19-22 October 2011

[arXiv:1203.5813]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

Focus beyond Quadratic Speedups for Error-Corrected Quantum Advantage

Ryan Babbush[*], Jarrod R. McClean,[†] Michael Newman, Craig Gidney, Sergio Boixo, and Hartmut Neven

Google Quantum AI, Venice, California 90291, USA

In this perspective we discuss conditions under which it would be possible for a modest fault-tolerant quantum computer to realize a runtime advantage by executing a quantum algorithm with only a small polynomial speedup over the best classical alternative. The challenge is that the computation must finish within a reasonable amount of time while being difficult enough that the small quantum scaling advantage would compensate for the large constant factor overheads associated with error correction. We compute several examples of such runtimes using state-of-the-art surface code constructions under a variety of assumptions. We conclude that quadratic speedups will not enable quantum advantage on early generations of such fault-tolerant devices unless there is a significant improvement in how we realize quantum error correction. While this conclusion persists even if we were to increase the rate of logical gates in the surface code by more than an order of magnitude, we also repeat this analysis for speedups by other polynomial degrees and find that quartic speedups look significantly more practical.

[arXiv:2011.04149]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

How Much Structure Is Needed for Huge Quantum Speedups?

Scott Aaronson*

September 2022

Abstract

I survey, for a general scientific audience, three decades of research into which sorts of problems admit exponential speedups via quantum computers—from the classics (like the algorithms of Simon and Shor), to the breakthrough of Yamakawa and Zhandry from April 2022. I discuss both the quantum circuit model, which is what we ultimately care about in practice but where our knowledge is radically incomplete, and the so-called *oracle* or *black-box* or *query complexity* model, where we've managed to achieve a much more thorough understanding that then informs our conjectures about the circuit model. I discuss the strengths and weaknesses of switching attention to sampling tasks, as was done in the recent quantum supremacy experiments. I make some skeptical remarks about widely-repeated claims of exponential quantum speedups for practical machine learning and optimization problems. Through many examples, I try to convey the "law of conservation of weirdness," according to which every problem admitting an exponential quantum speedup must have some unusual property to allow the amplitude to be concentrated on the unknown right answer(s).

Edited transcript of a rapporteur talk delivered at the 28th Solvay Physics Conference in Brussels, Belgium on May 21, 2022.

[arXiv:2209.06930]

⋆ What problems are "hard" on classical computers, which might be efficient for quantum computers?

Assessing requirements to scale to practical quantum advantage

M. E. Beverland,[1] P. Murali,[1] M. Troyer,[1] K. M. Svore,[1] T. Hoefler,[2]
V. Kliuchnikov,[1] G. H. Low,[1] M. Soeken,[3] A. Sundaram,[1] and A. Vaschillo[1]

[1]*Microsoft Quantum, Redmond, WA 98052, USA*
[2]*ETH Zurich, Department of Computer Science, Zürich, 8006, Switzerland*
[3]*Microsoft Quantum, Zurich, Switzerland*
(Dated: November 19, 2022)

While quantum computers promise to solve some scientifically and commercially valuable problems thought intractable for classical machines, delivering on this promise will require a *large-scale* quantum machine. Understanding the impact of architecture design choices for a scaled quantum stack for specific applications, prior to full realization of the quantum system, is an important open challenge. To this end, we develop a framework for quantum resource estimation, abstracting the layers of the stack, to estimate resources required across these layers for large-scale quantum applications. Using a tool that implements this framework, we assess three scaled quantum applications and find that hundreds of thousands to millions of physical qubits are needed to achieve practical quantum advantage. We identify three qubit parameters, namely size, speed, and controllability, that are critical at scale to rendering these applications practical. A goal of our work is to accelerate progress towards practical quantum advantage by enabling the broader community to explore design choices across the stack, from algorithms to qubits.

[arXiv:2211.07629]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

Disentangling Hype from Practicality: On Realistically Achieving Quantum Advantage

TORSTEN HOEFLER, Microsoft Corporation, USA and ETH Zurich, Switzerland
THOMAS HÄNER* and MATTHIAS TROYER, Microsoft Corporation, USA

Quantum computers offer a new paradigm of computing with the potential to vastly outperform any imagineable classical computer. This has caused a gold rush towards new quantum algorithms and hardware. In light of the growing expectations and hype surrounding quantum computing we ask the question which are the promising applications to realize quantum advantage. We argue that small data problems and quantum algorithms with super-quadratic speedups are essential to make quantum computers useful in practice. With these guidelines one can separate promising applications for quantum computing from those where classical solutions should be pursued. While most of the proposed quantum algorithms and applications do not achieve the necessary speedups to be considered practical, we already see a huge potential in material science and chemistry. We expect further applications to be developed based on our guidelines.

[arXiv:2307.00523]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

A framework for quantum advantage

O. Lanes,[1] M. Beji,[2] A. D. Córcoles,[1] C. Dalyac,[2] J. M. Gambetta,[1] L. Henriet,[2] A. Javadi-Abhari,[1] A. Kandala,[1] A. Mezzacapo,[1] C. Porter,[1] S. Sheldon,[1] J. Watrous,[1] C. Zoufal,[3] A. Dauphin,[2] and B. Peropadre[4]

[1] IBM Quantum, IBM Thomas J. Watson Research Center, USA[*]
[2] PASQAL SAS, 24 rue Emile Baudot - 91120 Palaiseau, Paris, France[†]
[3] IBM Quantum, IBM Research Europe – Zurich, 8803 Rüschlikon, Switzerland
[4] IBM Quantum, IBM Research, Cambridge, MA 02142, USA[‡]

As quantum computers approach the threshold where they can demonstrably outpace their classical counterparts on certain tasks, the need for a precise, consensus-driven definition of quantum advantage becomes essential. Rapid progress in the field has blurred this term across organizations, architectures, and application domains. Here, we aim to articulate a functional definition for quantum advantage that is both platform-agnostic and empirically verifiable. Building on this framework, we highlight the algorithmic families most likely to achieve early advantage. Finally, we outline our vision for the near future, in which quantum computers enhance existing high-performance computing platforms.

[arXiv:2506.20658]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

Formal Framework for Quantum Advantage

Harry Buhrman, Niklas Galke, Konstantinos Meichanetzidis

Quantinuum, Partnership House, Carlisle Place, London SW1P 1BX, United Kingdom
(Dated: October 3, 2025)

Motivated by notions of quantum heuristics and by average-case rather than worst-case algorithmic analysis, we define quantum computational advantage in terms of *individual* problem instances. Inspired by the classical notions of Kolmogorov complexity and instance complexity, we define their quantum versions. This allows us to define *queasy* instances of computational problems, like e.g. Satisfiability and Factoring, as those whose quantum instance complexity is significantly smaller than their classical instance complexity. These instances indicate quantum advantage: they are easy to solve on a quantum computer, but classical algorithms struggle (they feel queasy). Via a reduction from Factoring, we prove the existence of queasy Satisfiability instances; specifically, these instances are maximally queasy (under reasonable complexity-theoretic assumptions). Further, we show that there is exponential algorithmic utility in the queasiness of a quantum algorithm. This formal framework serves as a beacon that guides the hunt for quantum advantage in practice, and moreover, because its focus lies on single instances, it can lead to new ways of designing quantum algorithms.

[arXiv:2510.01953]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

Mind the gaps: The fraught road to quantum advantage

Jens Eisert[1,2,3] and John Preskill[4,5]

[1]*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*
[2]*Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany*
[3]*Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany*
[4]*Institute for Quantum Information and Matter, California Institute of Technology, CA 91125, USA*
[5]*AWS Center for Quantum Computing, Pasadena, CA 91125, USA*

Quantum computing is advancing rapidly, yet substantial gaps separate today's noisy intermediate-scale quantum (NISQ) devices from tomorrow's fault-tolerant application-scale quantum (FASQ) machines. We identify four related hurdles along the road ahead: (i) from error mitigation to active error detection and correction, (ii) from rudimentary error correction to scalable fault tolerance, (iii) from early heuristics to mature, verifiable algorithms, and (iv) from exploratory simulators to credible advantage in quantum simulation. Targeting these transitions will accelerate progress toward broadly useful quantum computing.

[arXiv:2510.19928]

★ What problems are "hard" on classical computers, which might be efficient for quantum computers?

The Grand Challenge of Quantum Applications

Ryan Babbush,[*] Robbie King,[†] Sergio Boixo, William Huggins, Tanuj Khattar,
Guang Hao Low, Jarrod R. McClean, Thomas O'Brien, and Nicholas C. Rubin

Google Quantum AI, Santa Barbara, CA 93111, United States
(Dated: December 5, 2025)

This perspective outlines promising pathways and critical obstacles on the road to developing useful quantum computing applications, drawing on insights from the Google Quantum AI team. We propose a five-stage framework for this process, spanning from theoretical explorations of quantum advantage to the practicalities of compilation and resource estimation. For each stage, we discuss key trends, milestones, and inherent scientific and sociological impediments. We argue that two central stages—identifying concrete problem instances expected to exhibit quantum advantage, and connecting such problems to real-world use cases—represent essential and currently under-resourced challenges. Throughout, we touch upon related topics, including the promise of generative artificial intelligence for aspects of this research, criteria for compelling demonstrations of quantum advantage, and the future of compilation as we enter the era of early fault-tolerant quantum computing.

[arXiv:2511.01924]

⋆ **What is the boundary between classical and quantum?**

(3) *No known classical algorithm can simulate a quantum computer.* But perhaps the most persuasive argument we have that quantum computing is powerful is simply that we don't know how to simulate a quantum computer using a digital computer; that remains true even after many decades of effort by physicists to find better ways to simulate quantum systems.

[arXiv:1801.00862]

Classical simulation

⋆ **What is the boundary between classical and quantum?**

(3) *No known classical algorithm can simulate a quantum computer.* But perhaps the most persuasive argument we have that quantum computing is powerful is simply that we don't know how to simulate a quantum computer using a digital computer; that remains true even after many decades of effort by physicists to find better ways to simulate quantum systems.

[arXiv:1801.00862]

▶ Roughly: we can simulate low-entangled states using tensor networks.

▶ Many approximation algorithms have been developed, e.g., in condensed matter theory and quantum chemistry.

▶ Sufficiently noisy circuits can be classically simulated!

★ What is the boundary between classical and quantum?

A polynomial-time classical algorithm for noisy random circuit sampling

Dorit Aharonov[*]   Xun Gao[†]   Zeph Landau[‡]   Yunchao Liu[§]   Umesh Vazirani[¶]

Abstract

We give a polynomial time classical algorithm for sampling from the output distribution of a noisy random quantum circuit in the regime of anti-concentration to within inverse polynomial total variation distance. This gives strong evidence that, in the presence of a constant rate of noise per gate, random circuit sampling (RCS) cannot be the basis of a scalable experimental violation of the extended Church-Turing thesis. Our algorithm is not practical in its current form, and does not address finite-size RCS based quantum supremacy experiments.

[arXiv:2211.03999]

⋆ **Quantum algorithms sometimes inspire better classical algorithms!**
(classical algorithms with only polynomial slowdown compared to the quantum one)

A quantum-inspired classical algorithm for recommendation systems

Ewin Tang

May 10, 2019

Abstract

We give a classical analogue to Kerenidis and Prakash's quantum recommendation system, previously believed to be one of the strongest candidates for provably exponential speedups in quantum machine learning. Our main result is an algorithm that, given an $m \times n$ matrix in a data structure supporting certain ℓ^2-norm sampling operations, outputs an ℓ^2-norm sample from a rank-k approximation of that matrix in time $O(\text{poly}(k)\log(mn))$, only polynomially slower than the quantum algorithm. As a consequence, Kerenidis and Prakash's algorithm does not in fact give an exponential speedup over classical algorithms. Further, under strong input assumptions, the classical recommendation system resulting from our algorithm produces recommendations exponentially faster than previous classical systems, which run in time linear in m and n.

The main insight of this work is the use of simple routines to manipulate ℓ^2-norm sampling distributions, which play the role of quantum superpositions in the classical setting. This correspondence indicates a potentially fruitful framework for formally comparing quantum machine learning algorithms to classical machine learning algorithms.

[arXiv:1807.04271]

★ **Quantum algorithms sometimes inspire better classical algorithms!**
(classical algorithms with only polynomial slowdown compared to the quantum one)

Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions

Ewin Tang[*]
University of Washington
(Dated: August 10, 2021)

A central roadblock to analyzing quantum algorithms on quantum states is the lack of a comparable input model for classical algorithms. Inspired by recent work of the author [E. Tang, STOC'19], we introduce such a model, where we assume we can efficiently perform ℓ^2-norm samples of input data, a natural analogue to quantum algorithms that assume efficient state preparation of classical data. Though this model produces less practical algorithms than the (stronger) standard model of classical computation, it captures versions of many of the features and nuances of quantum linear algebra algorithms. With this model, we describe classical analogues to Lloyd, Mohseni, and Rebentrost's quantum algorithms for principal component analysis [Nat. Phys. 10, 631 (2014)] and nearest-centroid clustering [arXiv:1307.0411]. Since they are only polynomially slower, these algorithms suggest that the exponential speedups of their quantum counterparts are simply an artifact of state preparation assumptions.

[arXiv:1811.00414]

⋆ **Quantum algorithms sometimes inspire better classical algorithms!**
(classical algorithms with only polynomial slowdown compared to the quantum one)

Revisiting dequantization and quantum advantage in learning tasks

Jordan Cotler,[1,2,*] Hsin-Yuan Huang,[3,4,†] and Jarrod R. McClean[5,‡]

[1]*Harvard Society of Fellows, Cambridge, MA 02138 USA*
[2]*Black Hole Initiative, Cambridge, MA 02138 USA*
[3]*Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA*
[4]*Department of Computing and Mathematical Sciences, Caltech, Pasadena, CA, USA*
[5]*Google Quantum AI, 340 Main Street, Venice, CA 90291, USA*
(Dated: December 7, 2021)

[arXiv:2112.00811]

Summary & goals of this class

▶ Quantum computers are known to have advantages over classical for certain tasks

▶ But realizing the advantage in practice is hard!

▶ Current quantum computers are too small and too noisy to yield advantage.

▶ There is an intense focus on building larger quantum computers with error correction. These devices may still be too small for the main applications.

▶ So there is intense focus on finding applications that can yield an advantage in the next ≈10 years with such small.

▶ But classical computational algorithms also improve at the same time!

Goals of the class

▶ Understand the basic language of quantum information and computation — this requires math!

▶ Learn the basic quantum algorithms (Shor's, Grover's, phase estimation).

▶ Understand the basics of how to correct errors in quantum computers.