
Quantum Computing: An Outlook for Platforms and Applications

Fred Chong

Seymour Goodman Professor, University of Chicago
Chief Scientist for Quantum Software, Infleqtion




THE UNIVERSITY OF
CHICAGO



Why Quantum Computing?



- Fundamentally change what is computable
 - The only means to potentially scale computation exponentially with the number of devices
- Solve currently intractable problems in chemistry, simulation, and optimization
 - Could lead to new nanoscale materials, better photovoltaics, better nitrogen fixation, and more



- A new industry and scaling curve to accelerate key applications
 - Not a full replacement for Moore's Law, but perhaps helps in key domains
- Lead to more insights in classical computing
 - Previous insights in chemistry, physics and cryptography
 - Challenge classical algorithms to compete w/ quantum algorithms

Why Now?

Now is a privileged time in the history of science and technology, as we are witnessing the opening of the NISQ era (where NISQ = noisy intermediate-scale quantum).


– John Preskill, Caltech



IBM
156 superconductor qubits



Inflektion
100+ neutral atom qubits




Quantinuum
98 atomic ion qubits

“Good” Quantum Applications

- Compact problem representation
 - Functions, small molecules, small graphs
- High complexity computation
- Compact solution
- Easily-verifiable solution
- Co-processing with classical supercomputers
- Can exploit a small number of quantum kernels



Quantum Bits (qubit)



- 1 qubit probabilistically represents 2 states


$$|a\rangle = C_0|0\rangle + C_1|1\rangle$$

- Every additional qubit doubles # states

$$|ab\rangle = C_{00}|00\rangle + C_{01}|01\rangle + \\ C_{10}|10\rangle + C_{11}|11\rangle$$


- *Quantum parallelism* on an exponential number of states
- But measurement collapses qubits to single classical values

Neutral Atom Quantum Computer




[Radnaev+ 2024]
(Inflection)

Quantum is Good for Quantum Simulation





Quantum Adiabatic Algorithm



Adiabatic evolution


Quantum Approximate Optimization Algorithm (QAOA)



Credit: <https://in.mathworks.com/help/matlab/math/max-cut-with-qaoa.html>




Hybrid Algorithms for Finding Multimodal Cancer Biomarkers



Co-Design Strategy

- Pick a classically difficult problem
 - ⑩ *Multimodal Cancer Biomarker Identification ($> 3^{\text{rd}}$ order correlations)*
 - ⑩ *Maybe so difficult that the classical algorithms still need to be developed*
- Preference for problems with smaller data requirements
 - ⑩ *Biomarker data acquisition is expensive ($\sim 10\text{K}$ records)*
 - ⑩ *Precision, rather than speed is more valuable (get the most out of valuable data)*
- Look for where classical methods break
 - ⑩ Apply quantum algorithms if possible

Why is Biomarker Discovery Challenging?



- Current biomarkers based on few features often **fail to provide clinical guidance**
 - Example: Many sequenced tumors do not have actionable mutations
- AI-based high-dimensional models can capture biological complexity but often lead to models that are **uninterpretable** and **fail to generalize, limiting clinical translation**, despite sophisticated computation

Key Challenges Summarized:

Sample Scarcity

Limited patient samples relative to feature space → overfitting, no generalization


Missing Synergies

Current methods miss higher-order interactions across modalities (eg DNA × RNA × Path)

Computational Limits

Classical solvers struggle with combinatorial complexity of feature selection at scale


How does the quantum computer help?



Bootstrapping from Smaller Problems


The first demonstrations of parameter transfer for RQAOA and for real-world problems

Variational optimization is *expensive*



- Many quantum circuit evaluations required to traverse the loss landscape
 - Limits the scalability of our classical simulations and hardware experiments


Because QAOA exhibits the *parameter concentration*¹ property, we can *transfer* parameters between problems



[1] Brandao, Fernando GSL, Michael Broughton, Edward Farhi, Sam Gutmann, and Hartmut Neven. "For fixed control parameters the quantum approximate optimization algorithm's objective function value concentrates for typical instances." *arXiv preprint arXiv:1812.04170* (2018).


Key Results Up Front

20-qubit QPU experiments demonstrate that our **hybrid algorithm boosts solver performance.**




Our **entropy-cubo**
(**H-CUBO**) feature selection
method outperforms
other baselines

Model (ERF) performance in predicting treatment response (pCR) in TN-like breast cancer patients (I-SPY2)




H-CUBO leverages **interactions across pathways** to predict treatment response






Executing 3 rounds of RQAOA on ibm_boston boosted the performance of QBSolv


TREX+ZNE+DD ZNE+DD
DD only No mitigation



- Maximize tractable circuit depth by layering different error mitigation techniques
- Shrink the effective problem size by running consecutive rounds of edge fixing




Input 20-variable problem






Reduced 17-variable problem

H-CUBO biomarkers better predict pCR in I-SPY2 patients than established biomarkers

Triple-negative (TN) breast cancer:
Difficult to treat



HER2+ breast cancer:
Easier to treat



H-CUBO biomarkers predict treatment outcome in important breast cancer subtypes

These biomarkers can be used to improve patient outcomes


Classical feature selection prioritizes "monergistic" features at the cost of synergy

CCL2 and *HSBP1* are in many synergistic triplets that FI does **not** leverage

HSBP1: negative regulator of heat shock response and cell stress, can also act as an "alarmin" to notify immune system of tissue damage

FI-selected genes include several other chemokines, likely important for directing immune cells to tumor microenvironment

Mutual Information (MI) with Treatment Response (pCR) of Triplets of Features in H-CUBO vs FI Selected Feature Sets



H-CUBO identifies subtle, synergistic interactions that connect biological pathways


Together with *CCL2*, H-CUBO solution includes *HSBP1*, *MYO5C*, *ABAT*, and *STAMBPL1*

This group highlights less well known, potentially impactful biology!




"HSPB1 facilitates chemoresistance through inhibiting ferroptotic cancer cell death and regulating NF- κ B signaling pathway in breast cancer" –Liang *et al.* (2023)

"Human myosin-Vc is a novel class V myosin expressed in epithelial cells" –Rodriguez & Cheney (2002)
Suggests MYO5C is involved in trafficking transferrin, a protein needed to bind/transport iron.

Mutual Information (MI) with Treatment Response (pCR) of Triplets of Features in H-CUBO vs FI Selected Feature Sets




Resource estimates in line with Starling Bicycle architecture



Starling
(100M)

Fault-tolerant

100M gates
200 logical qubits




Unlocking True Commercial Advantage

Designing data center ready solutions at scale


		2024 Error Detection	2026 Error Correction	2028 Commercial Advantage
Logical Performance		Dual-species crosstalk-free measurement	Logical circuit depth > 1,000 Universal quantum gate set	Logical circuit depth > 1million.
Logical Qubits		2	>10	>100
Logical Operations Rate			10,000/sec	100,000/sec
Physical qubits		1,600	8,000	40,000
Fidelity	2Q (CZ):	99.50%	99.90%	99.95%
	Local 1Q:	99.90%	99.95%	99.95%
	Global 1Q:	99.99%	99.99%	99.99%
Software		Tooling (ocvv) for verifying fault-tolerant properties	Optimized compilation of fault-tolerant circuits	Exponential speedup demonstration
Enabling technologies		Optimized laser pulse modulation	Advanced photonic beam steering	Realtime atom reloading

Demonstrating the power of full-stack co-design


- a) Fault tolerant decomposition to Clifford+T gateset can be prohibitively expensive.
- b) We can exploit the large design space of Algorithm+QEC Code to reduce hardware resource requirements.
- c) To illustrate this approach, we executed 12 logical qubit circuit on Inflektion Sqale using 32 physical qubits to sample low-energy solution to Entropy-CUBO.




(a) Solovay-Kitaev decomposition



(b) IQP circuits are easier to implement and they can sample low-energy distributions just like QAOA



(c) Encode within the $[[8,3,2]]$ error detecting code





Quantum-Inspired Contextual ML: Speedups with GPU and QPU



Enter: *Quantum Contextuality*



NOBELPRISET I FYSIK 2022
THE NOBEL PRIZE IN PHYSICS 2022




KUNGL.
VETENSKAPS-
AKADEMIEN
THE ROYAL SWEDISH ACADEMY OF SCIENCES




Alain Aspect

Université Paris-Saclay &
École Polytechnique, France



John F. Clauser

J.F. Clauser & Assoc.,
USA



Anton Zeilinger

University of Vienna,
Austria

*"för experiment med sammanflätade fotoner som påvisat brott mot Bell-olikheter och
banat väg för kvantinformationsvetenskap"*

*"for experiments with entangled photons, establishing the violation of Bell inequalities and
pioneering quantum information science"*

#nobelprize





Enter: *Quantum Contextuality*

ARTICLE

doi:10.1038/nature13460

Contextuality supplies the 'magic' for quantum computation

Mark Howard^{1,2}, Joel Wallman², Victor Veitch^{2,3} & Joseph Emerson²

Quantum computers promise dramatic advantages over their classical counterparts, but the source of the power in quantum computing has remained elusive. Here we prove a remarkable equivalence between the onset of contextuality and the possibility of universal quantum computation via 'magic state' distillation, which is the leading model for experimentally realizing a fault-tolerant quantum computer. This is a conceptually satisfying link, because it precludes a simple 'hidden variable' model of quantum mechanics, provides one of the fundamental uniquely quantum phenomena. Furthermore, this connection suggests a unifying paradigm for the reformation: the non-locality of quantum theory is a particular kind of contextuality, and non-locality to be a critical resource for achieving advantages with quantum communication. In addition to clarifying fundamental issues, this work advances the resource framework for quantum computation, which has a wide range of applications, such as characterizing the efficiency and trade-offs between distinct theoretical and experimental approaches to achieving robust quantum computation, and putting bounds on the overhead cost for the classical simulation of quantum algorithms.

Contextuality as a resource for measurement-based quantum computation beyond qubits

Markus Fremsbø¹, Sam Roberts² and Stephen D Bartlett¹

¹ Department of Physics, Imperial College London, London SW7 2AZ, United Kingdom

² Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, Australia

E-mail: stephen.bartlett@sydney.edu.au

Keywords: quantum information, quantum computing, contextuality, Bell nonlocality

Abstract

Contextuality—the obstruction to describing quantum mechanics in a classical statistical way—has been proposed as a resource that powers quantum computing. The measurement-based model provides a concrete manifestation of contextuality as a computational resource, as follows. If local measurements on a multi-qubit state can be used to evaluate nonlinear boolean functions with only linear control processing, then this computation constitutes a proof of strong contextuality—the possible local measurement outcomes cannot all be pre-assigned. However, this connection is restricted to the special case when the local measured systems are *qubits*, which have unusual properties from the perspective of contextuality. A single qubit cannot allow for a proof of contextuality, unlike higher-dimensional systems, and multiple qubits can allow for state-independent contextuality with only Pauli observables, again unlike higher-dimensional generalisations. Here we identify precisely that strong non-locality is necessary in a qudit measurement-based computation (MBC) that evaluates high-degree polynomial functions with only linear control. We introduce the concept of *local universality*, which places a bound on the space of output functions accessible under the constraint of single-qudit measurements. Thus, the partition of a physical system into subsystems plays a crucial role for the increase in computational power. A prominent feature of our setting is that the enabling resources for qubit and qudit MBC are of the same underlying nature, avoiding the pathologies associated with qubit contextuality.

Why are quantum algorithms so powerful?

Answer

Follow · 3

Request

3 Answers



Peter Shor · Following

Morss Professor of Applied Math at MIT. · Featured on Apple NewsUpvoted by Allan Steinhardt, PhD, Author "Radar in the Quantum Limit", Formerly DARPA's Chief Scientist, Fellow · 4y


They're more powerful because they are allowed to use quantum operations rather than classical operations, and quantum mechanics somehow has more computational power.

In the few years after Shor's factoring algorithm (that finds the prime factors of a large number) was discovered, there was a big controversy over which properties of quantum mechanics gave the extra computational power. Some people claimed that it was superposition, some people thought it was interference, some people believed it was the high-dimensionality of quantum state space that gave quantum computers their extra power. The controversy has now died down, and it appears that all of these people were right; you need all three of these properties of quantum mechanics to give quantum computers their extra power.



Contextual ML in a Nutshell

- *Quantum contextuality* enables wider “context clues” for AI/ML.
 - Ex: Forecasting. Energy demand prediction depends not only on previous minute, but also 24 hours ago, 7 days ago, 365 days ago, ...




Details: [1] Anschuetz, Eric, Hu, H.Y., Huang, J.L. and Gao, X., 2023. Interpretable quantum advantage in neural sequence learning. *PRX Quantum*, 4(2), p.020338. [2] Gao, X., Anschuetz, Eric, Wang, S.T., Cirac, J.I. and Lukin, M.D., 2022. Enhancing generative models via quantum correlations. *Physical Review X*, 12(2), p.021037.



Advantages of Contextual ML

1. **Deployment-ready.** Runs today on GPU, scales gracefully to *QPU*.




2. **Interpretability.** HMM-based approach has high interpretability. Important given regulatory requirements for fairness and safety.
3. **Per-dollar or per-watt performance.** Given a fixed budget, contextual ML achieves greater accuracy.

What's the catch? (a) Only useful for sequential data like time series.
(b) Lower absolute performance than GPT.



Contextual ML: QPU $\leftrightarrow$ GPU


- Key insight: better to store memory (hidden state) as a square matrix vs. a column matrix



Memory size is

Inference time scales as n^2

Number of weights is n^2



Memory size is n^2

Inference time scales as n^{ω} :


- 2.81 w/ Strassen
- 2.37 $\tilde{w}$ / best asymptotic

Number of weights is n^2



Timeseries Energy Data

- Dataset: Hourly energy consumption



Advanced Research Projects Agency • ENERGY




- Task: predict next hour of energy usage based on history
- Memory is the limiting factor in modern approaches.



Signal Understanding: Quantum-Inspired Rapid Context



- Navy contract leveraging CML-based autoencoders for signal understanding → effectively acts as compression for long signals.





Genomic Benchmarks

- Exemplar of context-sensitivity

DATABASE

Open Access

Genomic benchmarks: a collection
of datasets for genomic sequence classification



Katarína Grešová^{1,2}, Vlastimil Martinek^{1,2}, David Čechák^{1,2}, Petr Šimeček^{1*}  and Panagiotis Alexiou¹

```
>>> from genomic_benchmarks.dataset_getters.pytorch_datasets import get_dataset
>>>
>>> dset = get_dataset('human_nontata_promoters', split='train', version=0)
>>> dset[0]
('CAATCTCACAGGCTCCTGGTTGTCTACCCATGGACCCAGAGGTTCTTTGACAGCTTTGGCA...TCCAGGAGATGT', 0)
```

- Our focus: mouse enhancers ensemble → longest median length



Genomic Benchmarks

Model	Accuracy	Pre-training?	Size
CNN	69.0%	N	3 layers with 16, 8, 4 filters
DNABERT	66.9%	Y	110M params
Transformer	80.1% / 80.1%	Y / N (same accuracy)	530k params
HyenaDNA	85.1% / 84.7%	Y / N	440k params

github.com/ML-Bioinfo-CEITEC/genomic_benchmarks

README Apache-2.0 license

Genomic Benchmarks

In this repository, we collect benchmarks for classification of genomic sequences. It is shipped as a Python package, together with functions helping to download & manipulate datasets and train NN models. [Current SOTA model on genomic benchmarks is HyenaDNA](#), see metrics in the [experiments](#) folder.

HyenaDNA: Long-Range Genomic Sequence Modeling at Single Nucleotide Resolution

Eric Nguyen¹, Michael Poli^{1,2,*}, Marjan Faizi^{2,*},
Armin W. Thomas¹, Callum Birch Sykes³, Michael Wornow¹, Aman Patel¹,
Clayton Rabideau³, Stefano Massaroli⁴, Yoshua Bengio⁴, Stefano Ermon¹,
Stephen A. Baccus^{1,†}, Christopher Ré^{1,†}



Genomic Benchmarks

Model	Accuracy	Pre-training?	Size
CNN	69.0%	N	3 layers with 16, 8, 4 filters
DNABERT	66.9%	Y	110M params
Transformer	80.1% / 80.1%	Y / N (same accuracy)	530k params
HyenaDNA	85.1% / 84.7%	Y / N	440k params
CML	86.7%	N	240k params

A new record!



Contextual ML Summary

- Quantum contextuality yields:
 - ❑ Today, with GPU by reshaping memory
 - Enables larger memories within same inference time
 - ❑ Tomorrow, with QPU: linear runtime
- Inductive bias helps for many datasets including sensor data streams and energy forecasting.
- GPU (PyTorch) and QPU (CUDA-Q) implementations ready today

Quantum Sensing

■ Application Objective:

- ❑ MHz-GHz spectral sensing (QRF) for monitoring bioreactor health.
- ❑ Monitor metabolic processes, detect biomolecular interactions.
- ❑ Needed for early warning, process optimization, sterility assurance, quality control.

■ Current solution:

- ❑ Combination of spectroscopy, sampling, classical sensors.
- ❑ Limited by delayed feedback, intrusiveness, limited sensitivity, narrowband (kHz to low GHz)


■ Novel approach:

- ❑ Passive, wideband spectral analysis.
- ❑ Higher sensitivity to weak signals.
- ❑ Non-invasive.
- ❑ Signal interpretation with CML to extract meaningful features from overlapping signals.
- ❑ **NIH paper, show that QRF is sensitive to rotational states of molecules and biomarkers.**





Quantum Software: Please break abstraction layers!




X. Fu et al. (2017)
arXiv:1708.07677

Shortcut


- ▷ Stack: rigid layers + interfaces
- ▷ Benefits
 - Taming complexity
- ▷ **Problems**
 - Lost opportunities for optimization
 - QC stack + layers change
- ▷ We should compile to hardware primitives. Physics first.



Deep Optimization Extends the Computation Frontier



Hybrid Quantum-Classical Computing: Contest of Exponentials







Summary

- QC is at a historic time
- Hybrid quantum-classical compute will be key
- Quantum-inspired approaches can be useful in the short term
- Physics-aware, full-stack SW can greatly accelerate progress
- More info: epiqc.cs.uchicago.edu
inflection.com



BACKUP SLIDES

		Qubit Coherence Time (sec)	Two-qubit Gate Fidelity	Qubits Connected	Companies	Pros	Cons
Natural qubits							
	Trapped Ions Electrically charged atoms, or ions, are held in place with electric fields. Qubits are stored in electronic states. Ions are pushed with laser beams to allow the qubits to interact.	>1000	99.9%	High	IonQ, Quantinuum, AQFT Oxford Ionics, Universal Quantum	Very stable. Highest achieved gate fidelities.	Slow operation. Many lasers are needed.
	Neutral Atoms Neutral atoms, like ions, store qubits within electronic states. Laser activates the electrons to create interaction between qubits.	1	99.5%	Very high; low individual control	Infleqion, Atom Computing, QuEra, Pasqal, Ruanq, MP	Many qubits, 2D and maybe 3D.	Hard to program and control individual qubits; prone to noise.
	Photonics Photonic qubits are sent through a maze of optical channels on a chip to interact. At the end of the maze, the distribution of photons is measured as output.	-	-	-	PsiQuantum, Xanadu	Linear optical gates, integrated on-chip.	Each program requires its own chip with unique optical channels. No memory.
	Diamond Vacancies A nitrogen atom and a vacancy add an electron to a diamond lattice. Its quantum spin state, along with those of nearby carbon nuclei, can be controlled with light.	10	99.2%	Low	Quantum Diamond Technologies, Quantum Brilliance	Can operate at room temperature.	Difficult to create high numbers of qubits, limiting compute capacity.
Synthetic qubits							
	Superconducting Circuits A resistance-free current oscillates back and forth around a circuit loop. An injected microwave signal excites the current into super-position states.	0.00005	99.9%	High	Google, IBM, OQI, Rigetti, Oxford Quantum Circuits	Can lay out physical circuits on chip.	Must be cooled to near absolute zero. High variability in fabrication. Lots of noise.
	Silicon Quantum Dots These "artificial atoms" are made by adding an electron to a small piece of pure silicon. Microwaves control the electron's quantum state.	0.03	<99%	Very Low	HLRL, Intel, SQC, Oxford Quantum, Ocean, DRAQ, Quantum Motion, ZeroQ	Borrows from existing semiconductor industry.	Only a few connected. Must be cooled to near absolute zero. High variability in fabrication.
	Topological Qubits Quasiparticles can be seen in the behavior of electrons channeled through semiconductor structures. Their braided paths can encode quantum information.	-	-	-	Microsoft	Designed to be more robust to environmental noise.	Existence not yet confirmed.

5 Year Update to the Next Steps in Quantum Computing Report