Distributed quantum machine learning under practical communication schemes

Yosep Kim

Department of Physics,
Korea University

Superconducting
Quantum Information Lab

Introduction

Quantum supremacy?

In **2019**

Article

Quantum supremacy using a programmable superconducting processor

https://doi.org/10.1038/s41586-019-1666-5
Received: 22 July 2019
Accepted: 20 September 2019
Published online: 23 October 2019

Frank Arute[1], Kunal Arya[1], Ryan Babbush[1], Dave Bacon[1], Joseph C. Bardin[1,2], Rami Barends[1], Rupak Biswas[3], Sergio Boixo[1], Fernando G. S. L. Brandao[1,4], David A. Buell[1], Brian Burkett[1], Yu Chen[1], Zijun Chen[1], Ben Chiaro[5], Roberto Collins[1], William Courtney[1], Andrew Dunsworth[1], Edward Farhi[1], Brooks Foxen[1,5], Austin Fowler[1], Craig Gidney[1], Marissa Giustina[1], Rob Graff[1], Keith Guerin[1], Steve Habegger[1], Matthew P. Harrigan[1], Michael J. Hartmann[1,6], Alan Ho[1], Markus Hoffmann[1], Trent Huang[1], Travis S. Humble[7], Sergei V. Isakov[1], Evan Jeffrey[1], Zhang Jiang[1], Dvir Kafri[1], Kostyantyn Kechedzhi[1], Julian Kelly[1], Paul V. Klimov[1], Sergey Knysh[1], Alexander Korotkov[1,8], Fedor Kostritsa[1], David Landhuis[1], Mike Lindmark[1], Erik Lucero[1], Dmitry Lyakh[9], Salvatore Mandrà[3,10], Jarrod R. McClean[1], Matthew McEwen[1,5], Anthony Megrant[1], Xiao Mi[1], Kristel Michielsen[11,12], Masoud Mohseni[1], Josh Mutus[1], Ofer Naaman[1], Matthew Neeley[1], Charles Neill[1], Murphy Yuezhen Niu[1], Eric Ostby[1], Andre Petukhov[1], John C. Platt[1], Chris Quintana[1], Eleanor G. Rieffel[3], Pedram Roushan[1], Nicholas C. Rubin[1], Daniel Sank[1], Kevin J. Satzinger[1], Vadim Smelyanskiy[1], Kevin J. Sung[1,13], Matthew D. Trevithick[1], Amit Vainsencher[1], Benjamin Villalonga[1,14], Theodore White[1], Z. Jamie Yao[1], Ping Yeh[1], Adam Zalcman[1], Hartmut Neven[1] & John M. Martinis[1,5]*

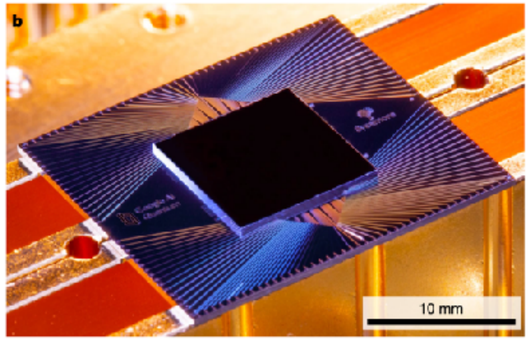

Google AI Quantum

Our Sycamore processor takes about **200 seconds** to sample one instance of a quantum circuit a million times—our benchmarks currently indicate that the equivalent task for a state-of-the-art classical supercomputer would take approximately **10,000 years**.

In **2023**

Article

Evidence for the utility of quantum computing before fault tolerance

https://doi.org/10.1038/s41586-023-06096-3
Received: 24 February 2023
Accepted: 18 April 2023
Published online: 14 June 2023

Youngseok Kim[1,6]✉, Andrew Eddins[2,6]✉, Sajant Anand[3], Ken Xuan Wei[1], Ewout van den Berg[1], Sami Rosenblatt[1], Hasan Nayfeh[1], Yantao Wu[3,4], Michael Zaletel[3,5], Kristan Temme[1] & Abhinav Kandala[1]✉

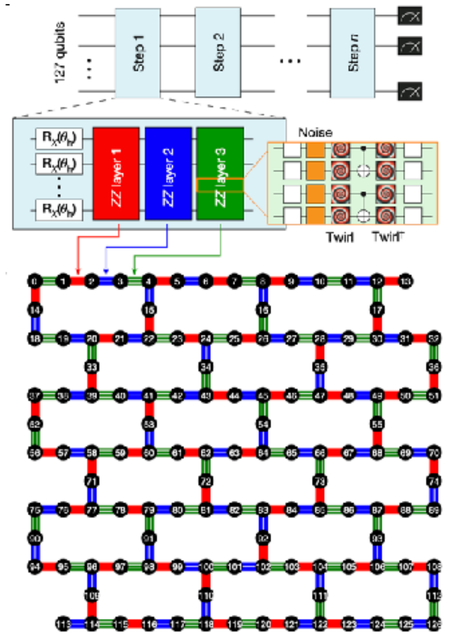

IBM **Quantum**

The tensor network simulations in Fig.4 were run on a 64-core, 2.45-GHz processor with 128 GB of memory, … **8 h** for Fig. 4a and **30 h** for Fig. 4b.

The corresponding quantum wall-clock run time was approximately **4 h** for Fig.4a and **9.5 h** for Fig.4b

Quantum supremacy?

In **2022**

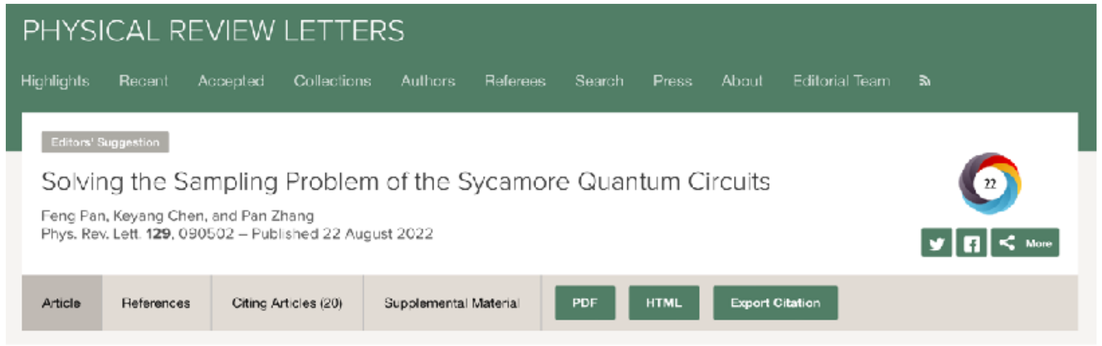

It would take only **a few dozen seconds** if they had access to full-sized supercomputers.

In **2023**

Fast classical simulation of evidence for the utility of quantum computing before fault tolerance

Tomislav Begušić and Garnet Kin-Lic Chan*
Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, California 91125, USA
(Dated: June 29, 2023)

We show that a classical algorithm based on sparse Pauli dynamics can efficiently simulate quantum circuits studied in a recent experiment on 127 qubits of IBM's Eagle processor [*Nature* **618**, 500 (2023)]. Our classical simulations on a single core of a laptop are orders of magnitude faster than the reported walltime of the quantum simulations, as well as faster than the estimated quantum hardware runtime without classical processing, and are in good agreement with the zero-noise extrapolated experimental results.

Our classical simulations on a single core of a laptop are **orders of magnitude faster** than the reported walltime of the quantum simulations

Quantum supremacy!

Article

Phase transitions in random circuit sampling

https://doi.org/10.1038/s41586-024-07998-6
Received: 9 January 2024
Accepted: 26 August 2024
Published online: 9 October 2024
Open access
Check for updates

Undesired coupling to the surrounding environment destroys long-range correlations in quantum processors and hinders coherent evolution in the nominally available computational space. This noise is an outstanding challenge when leveraging the computation power of near-term quantum processors[1]. It has been shown that benchmarking random circuit sampling with cross-entropy benchmarking can provide an estimate of the effective size of the Hilbert space coherently available[2–8]. Nevertheless, quantum algorithms' outputs can be trivialized by noise, making them susceptible to classical computation spoofing. Here, by implementing an algorithm for random circuit sampling, we demonstrate experimentally that two phase transitions are observable with cross-entropy benchmarking, which we explain theoretically with a statistical model. The first is a dynamical transition as a function of the number of cycles and is the continuation of the anti-concentration point in the noiseless case. The second is a quantum phase transition controlled by the error per cycle; to identify it analytically and experimentally, we create a weak-link model, which allows us to vary the strength of the noise versus coherent evolution. Furthermore, by presenting a random circuit sampling experiment in the weak-noise phase with 67 qubits at 32 cycles, we demonstrate that the computational cost of our experiment is beyond the capabilities of existing classical supercomputers. Our experimental and theoretical work establishes the existence of transitions to a stable, computationally complex phase that is reachable with current quantum processors.

Table 1 | Estimated computational cost of simulation

Experiment	*d*	One amplitude	1 million noisy samples		
		(FLOPs)	FLOPs	XEB fidelity	Time
SYC-53 (ref. 4)	20	5×10^{18}	2×10^{18}	2×10^{-3}	6 s
ZCZ-56 (ref. 5)	20	1×10^{20}	7×10^{20}	6×10^{-4}	20 min
ZCZ-60 (ref. 6)	24	6×10^{21}	3×10^{23}	3×10^{-4}	40 days
SYC-70	24	4×10^{24}	5×10^{26}	2×10^{-3}	50 years
SYC-67	32	1×10^{24}	1×10^{38}	1×10^{-3}	1×10^{13} years
			1×10^{29}		1×10^{4} years[a]
			1×10^{26}		12 years[b]

[a]We include the cost estimated by assuming memory is distributed over all RAM, ignoring realistic bandwidth constraints.
[b]We include the cost estimated by assuming memory is distributed over all secondary storage, ignoring realistic bandwidth constraints.

Quantum supremacy!

Random circuit sampling

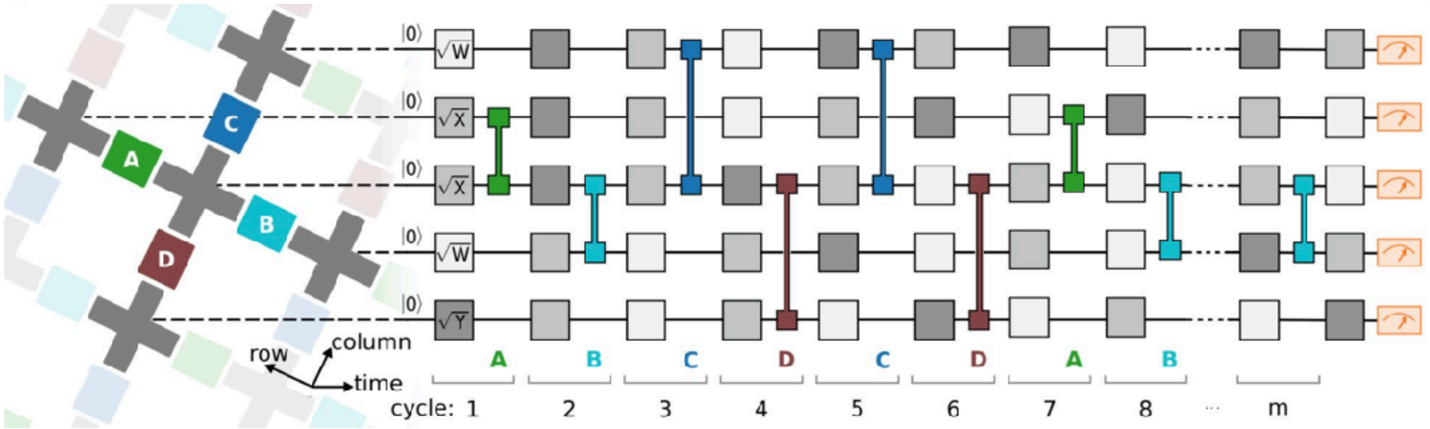

Unfortunately, such sampling problems are not very practical.

Quantum supremacy for practical problems?!

Article

Observation of constructive interference at the edge of quantum ergodicity

https://doi.org/10.1038/s41586-025-09526-6
Received: 3 November 2024
Accepted: 13 August 2025
Open access
Check for updates

Google Quantum AI and Collaborators*

The dynamics of quantum many-body systems is characterized by quantum observables that are reconstructed from correlation functions at separate points in space and time[1–3]. In dynamics with fast entanglement generation, however, quantum observables generally become insensitive to the details of the underlying dynamics at long times due to the effects of scrambling. To circumvent this limitation and enable access to relevant dynamics in experimental systems, repeated time-reversal protocols have been successfully implemented[4]. Here we experimentally measure the second-order out-of-time-order correlators ($OTOC^{(2)}$)[4–18] on a superconducting quantum processor and find that they remain sensitive to the underlying dynamics at long timescales. Furthermore, $OTOC^{(2)}$ manifests quantum correlations in a highly entangled quantum many-body system that are inaccessible without time-reversal techniques. This is demonstrated through an experimental protocol that randomizes the phases of Pauli strings in the Heisenberg picture by inserting Pauli operators during quantum evolution. The measured values of $OTOC^{(2)}$ are substantially changed by the protocol, thereby revealing constructive interference between Pauli strings that form large loops in the configuration space. The observed interference mechanism also endows $OTOC^{(2)}$ with high degrees of classical simulation complexity. These results, combined with the capability of $OTOC^{(2)}$ in unravelling useful details of quantum dynamics, as shown through an example of Hamiltonian learning, indicate a viable path to practical quantum advantage.

same circuits. Figure 4c shows the estimated cost of simulating $C^{(4)}_{\text{diag}}$ through tensor-network contraction on the Frontier supercomputer, which converges to approximately 3.2 years. This is a factor of approximately 13,000 longer than the experimental data collection time of 2.1 h per circuit, indicating that this experiment is, at present, in the beyond-classical regime of quantum computation.

Quantum computation of molecular geometry via many-body nuclear spin echoes

arXiv:2510.19550v1

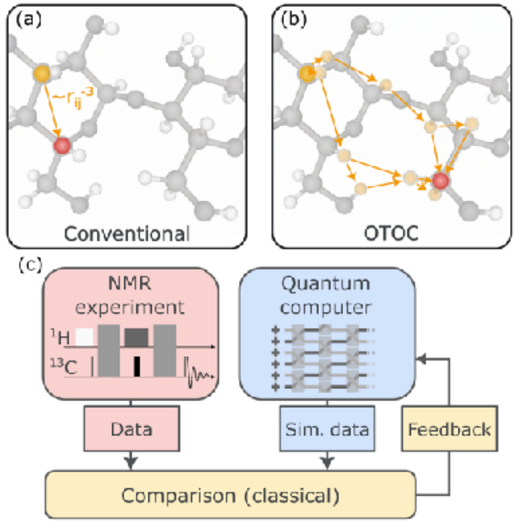

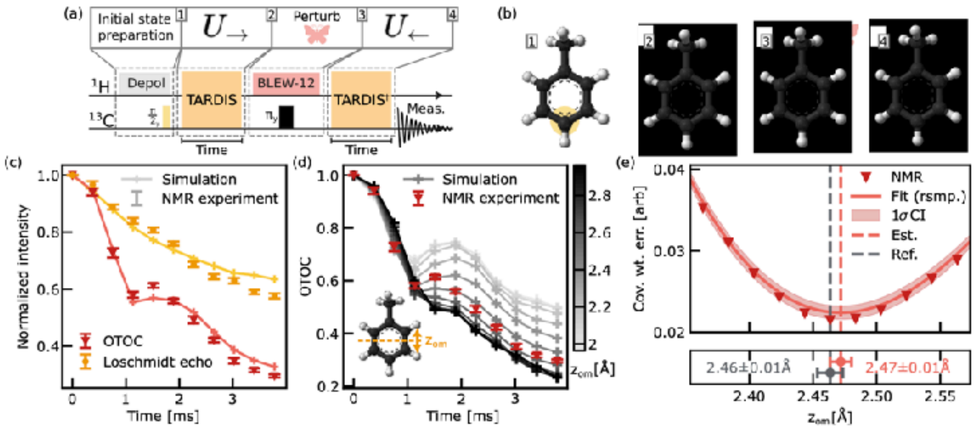

Noisy Intermediate-Scale Quantum (NISQ)

(a) Complexity class

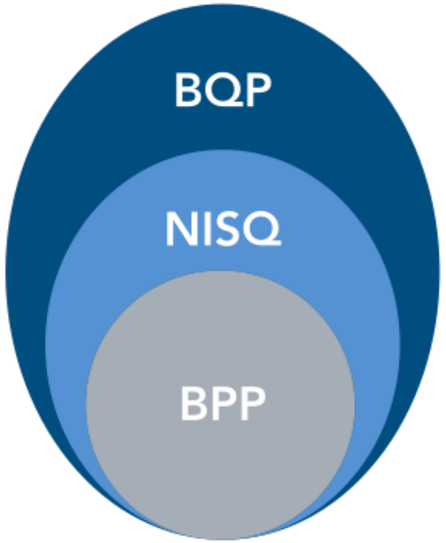

(b) An algorithm in NISQ

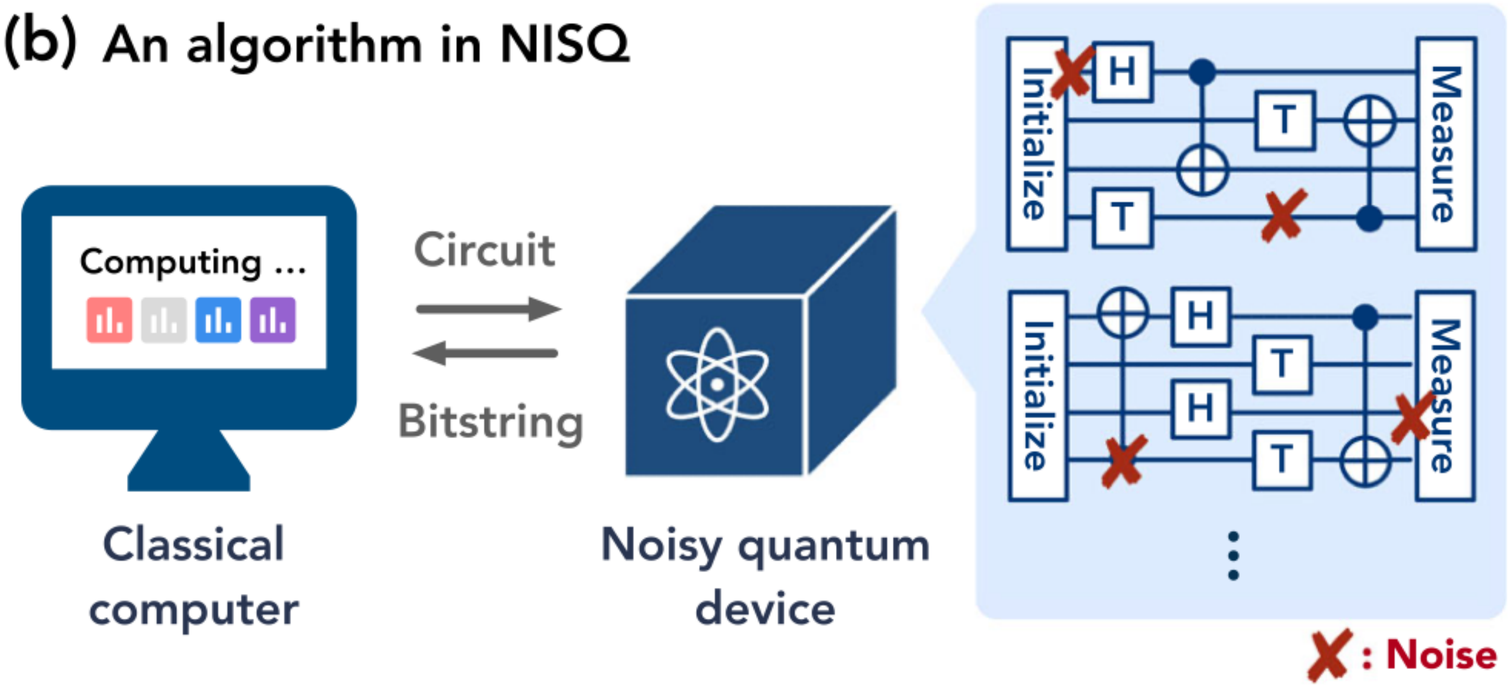

$$\mathrm{BPP} \subseteq \mathrm{NISQ} \subseteq \mathrm{BQP}$$

S. Chen, J. Cotler, H.-Y. Huang, J. Li, Nat. Commun. 14, 6001 (2023)

Development Roadmap

IBM Quantum

	2016–2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2033+
	Run quantum circuits on the IBM Quantum Platform	Release multi-dimensional roadmap publicly with initial aim focused on scaling	Enhancing quantum execution speed by 100x with Qiskit Runtime	Bring dynamic circuits to unlock more computations	Enhancing quantum execution speed by 5x with quantum serverless and Execution modes	Improving quantum circuit quality and speed to allow 5K gates with parametric circuits	Enhancing quantum execution speed and parallelization with partitioning and quantum modularity	Improving quantum circuit quality to allow 7.5K gates	Improving quantum circuit quality to allow 10K gates	Improving quantum circuit quality to allow 15K gates	Improving quantum circuit quality to allow 100M gates	Beyond 2033, quantum-centric supercomputers will include 1000's of logical qubits unlocking the full power of quantum computing
Data Scientist						Platform						
						Code assistant	Functions	Mapping Collection	Specific Libraries			General purpose QC libraries
Researchers					Middleware							
					Quantum Serverless	Transpiler Service	Resource Management	Circuit Knitting x P	Intelligent Orchestration			Circuit libraries
Quantum Physicist			Qiskit Runtime									
	IBM Quantum Experience		QASM3	Dynamic circuits	Execution Modes	Heron (5K)	Flamingo (5K)	Flamingo (7.5K)	Flamingo (10K)	Flamingo (15K)	Starling (100M)	Blue Jay (1B)
						Error Mitigation	Error Mitigation	Error Mitigation	Error Mitigation	Error Mitigation	Error correction	Error correction
						5k gates 133 qubits	5k gates 156 qubits	7.5k gates 156 qubits	10k gates 156 qubits	15k gates 156 qubits	100M gates 200 qubits	1B gates 2000 qubits
						Classical modular	Quantum modular	Quantum modular	Quantum modular	Quantum modular	Error corrected modularity	Error corrected modularity
						133x3 = 399 qubits	156x7 = 1092 qubits	156x7 = 1092 qubits	156x7 = 1092 qubits	156x7 = 1092 qubits		
	Early Canary 5 qubits Albatross 16 qubits Penguin 20 qubits Prototype 53 qubits	Falcon Benchmarking 27 qubits		Eagle Benchmarking 127 qubits								

Innovation Roadmap

Software Innovation	IBM Quantum Experience	Qiskit Circuit and operator API with compilation to multiple targets	Application modules Modules for domain specific application and algorithm workflows	Qiskit Runtime Performance and abstract through Primitives	Serverless Demonstrate concepts of quantum centric-supercomputing	AI enhanced quantum Prototype demonstrations of AI enhanced circuit transpilation	Resource management System partitioning to enable parallel execution	Scalable circuit knitting Circuit partitioning with classical reconstruction at HPC scale	Error correction decoder Demonstration of a quantum system with real-time error correction decoder			
Hardware Innovation	Early Canary 5 qubits Penguin 20 qubits Albatross 16 qubits Prototype 53 qubits	Falcon Demonstrate scaling with I/O routing with Bump bonds	Hummingbird Demonstrate scaling with multiplexing readout	Eagle Demonstrate scaling with MLW and TSV	Osprey Enabling scaling with high density signal delivery	Condor Single system scaling and fridge capacity	Flamingo Demonstrate scaling with modular connectors	Kookaburra Demonstrate scaling with nonlocal c-coupler	Demonstrate path to improved quality with logical memory	Cockatoo Demonstrate path to improved quality with logical communication	Starling Demonstrate path to improved quality with logical gates	
						Heron Architecture based on tunable-couplers	Crossbill m- coupler					

Executed by IBM

On target

	Execution Modes	Heron (5K)	Flamingo (5K)	Flamingo (7.5K)	Flamingo (10K)	Flamingo (15K)	Starling (100M)	Blue Jay (1B)
	2	Error Mitigation	Error Mitigation	Error Mitigation	Error Mitigation	Error Mitigation	Error correction	Error correction
		5k gates 133 qubits	5k gates 156 qubits	7.5k gates 156 qubits	10k gates 156 qubits	15k gates 156 qubits	100M gates 200 qubits	1B gates 2000 qubits
		Classical modular	Quantum modular	Quantum modular	Quantum modular	Quantum modular	Error corrected modularity	Error corrected modularity
		133x3 = 399 qubits	156x7 = 1092 qubits	156x7 = 1092 qubits	156x7 = 1092 qubits	156x7 = 1092 qubits		
	AI enhanced quantum	Resource management	Scalable circuit knitting	Error correction decoder				
	Prototype demonstrations of AI enhanced circuit transpilation	System partitioning to enable parallel execution	Circuit partitioning with classical reconstruction at HPC scale	Demonstration of a quantum system with real-time error correction decoder				
	Condor	Flamingo	Kookaburra		Cockatoo	Starling		
	Single system scaling and fridge capacity	Demonstrate scaling with modular connectors	Demonstrate scaling with nonlocal c-coupler	Demonstrate path to improved quality with logical memory	Demonstrate path to improved quality with logical communication	Demonstrate path to improved quality with logical gates		
	Heron	Crossbill						
	Architecture based on tunable-couplers	m- coupler						

Challenges in scaling
the qubit count per chip?!

Frequency allocation in multi-qubit processors

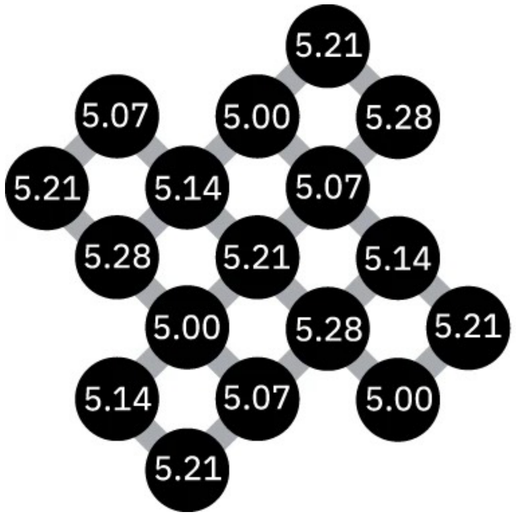

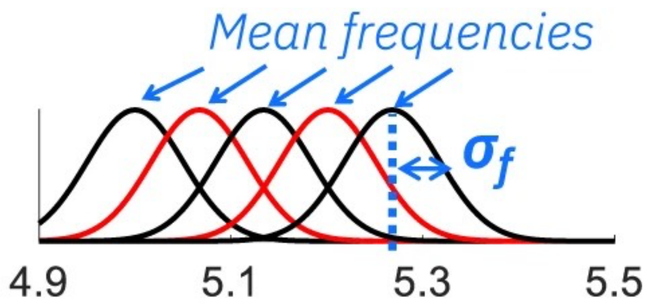

Frequency-crowding problem

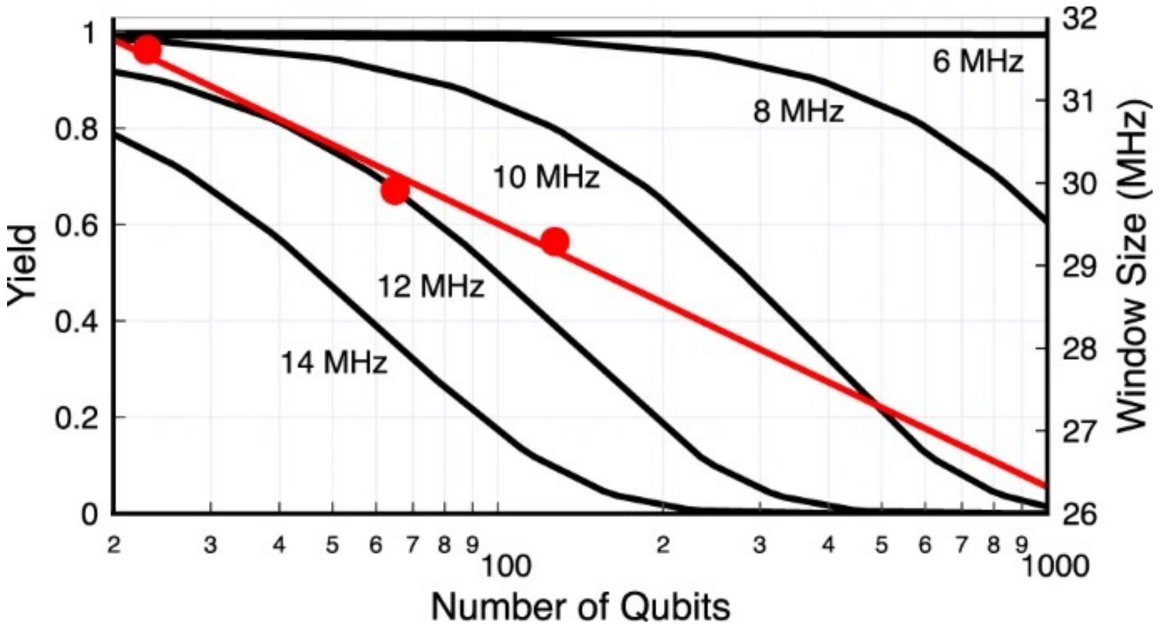

J. B. Hertzberg et al., npj Quantum Info. 7, 129 (2021)

Modular architecture

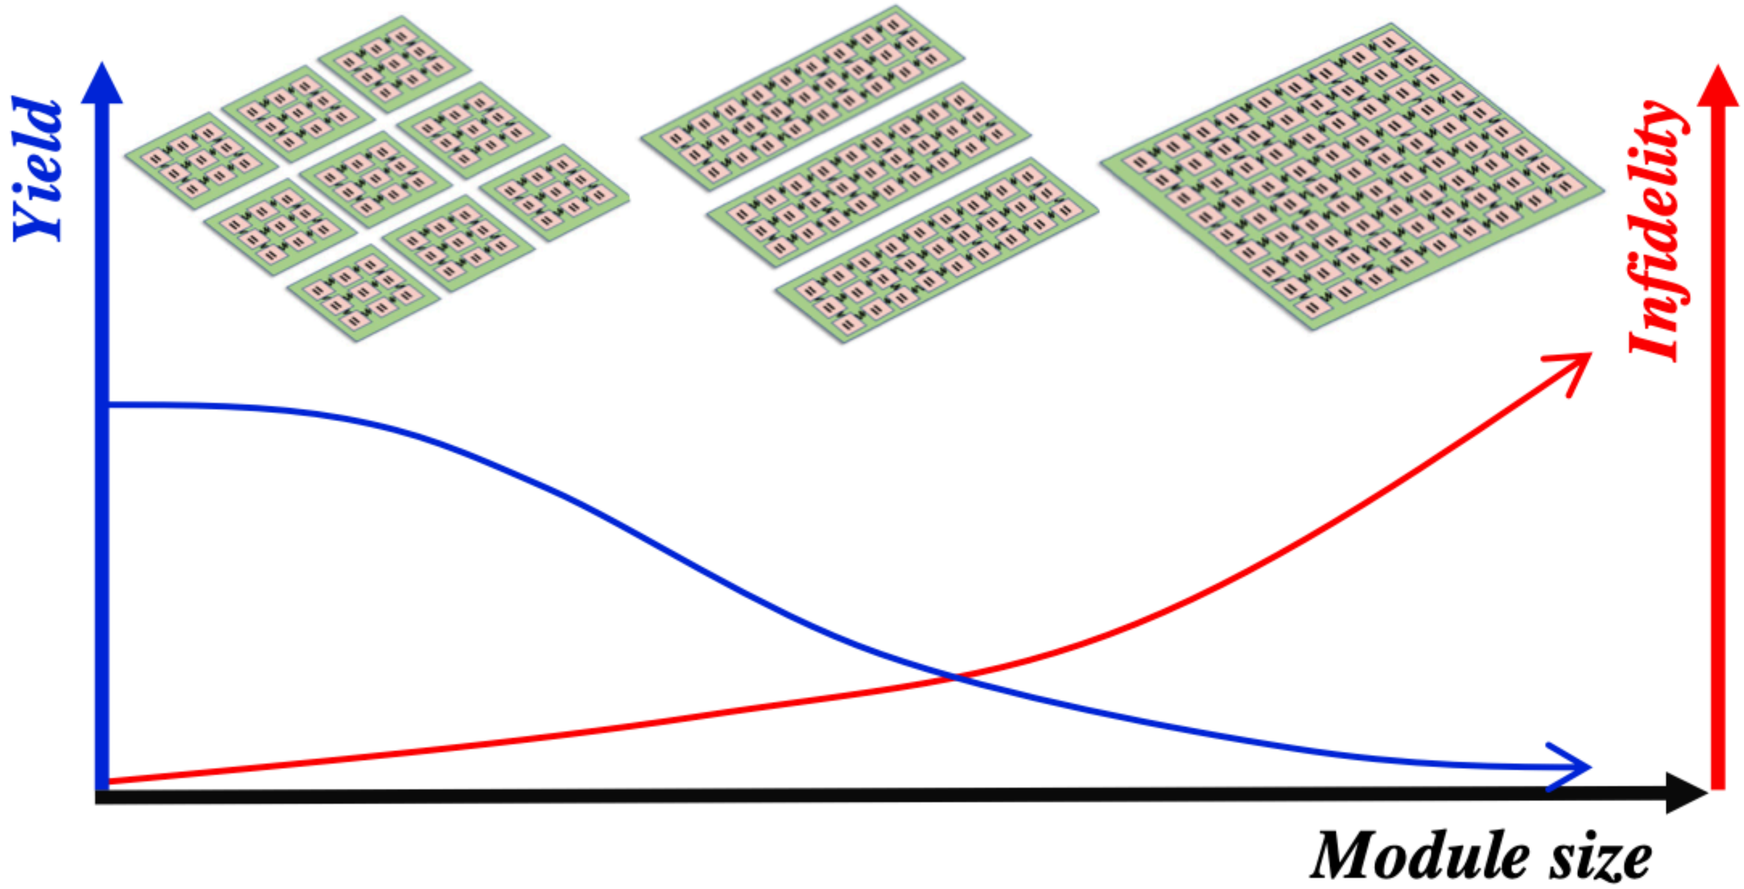

K. N. Smith et al., arXiv:2210.10921

Modular architecture

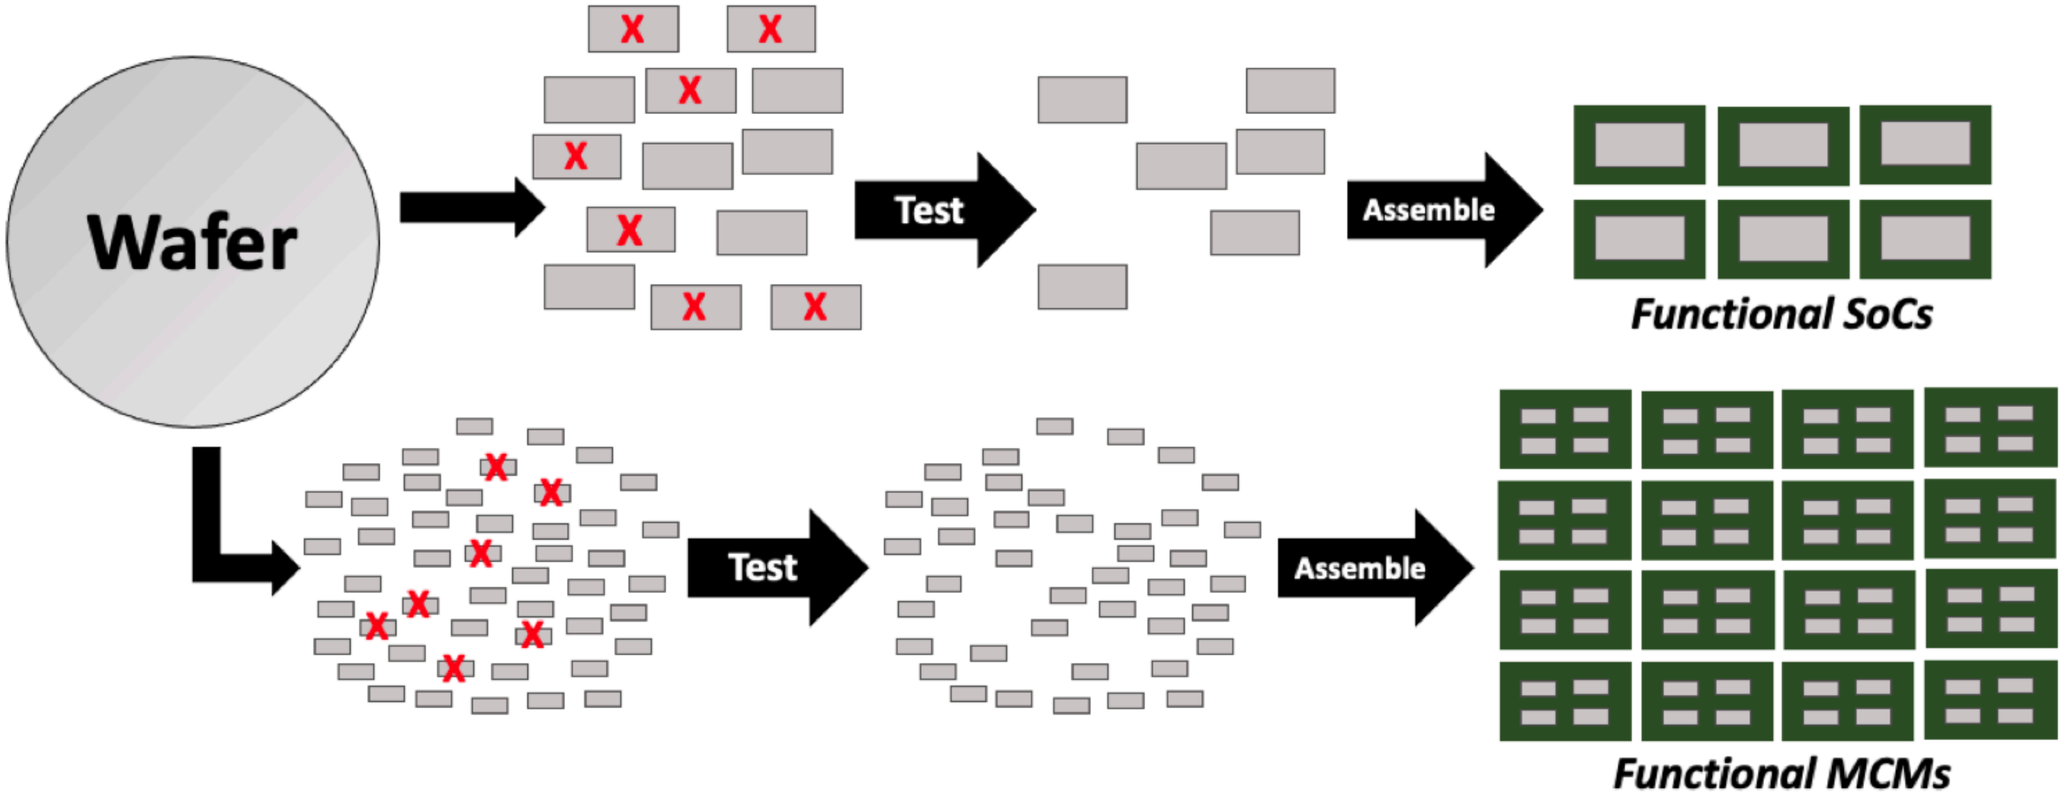

K. N. Smith et al., arXiv:2210.10921

Chiplet architecture

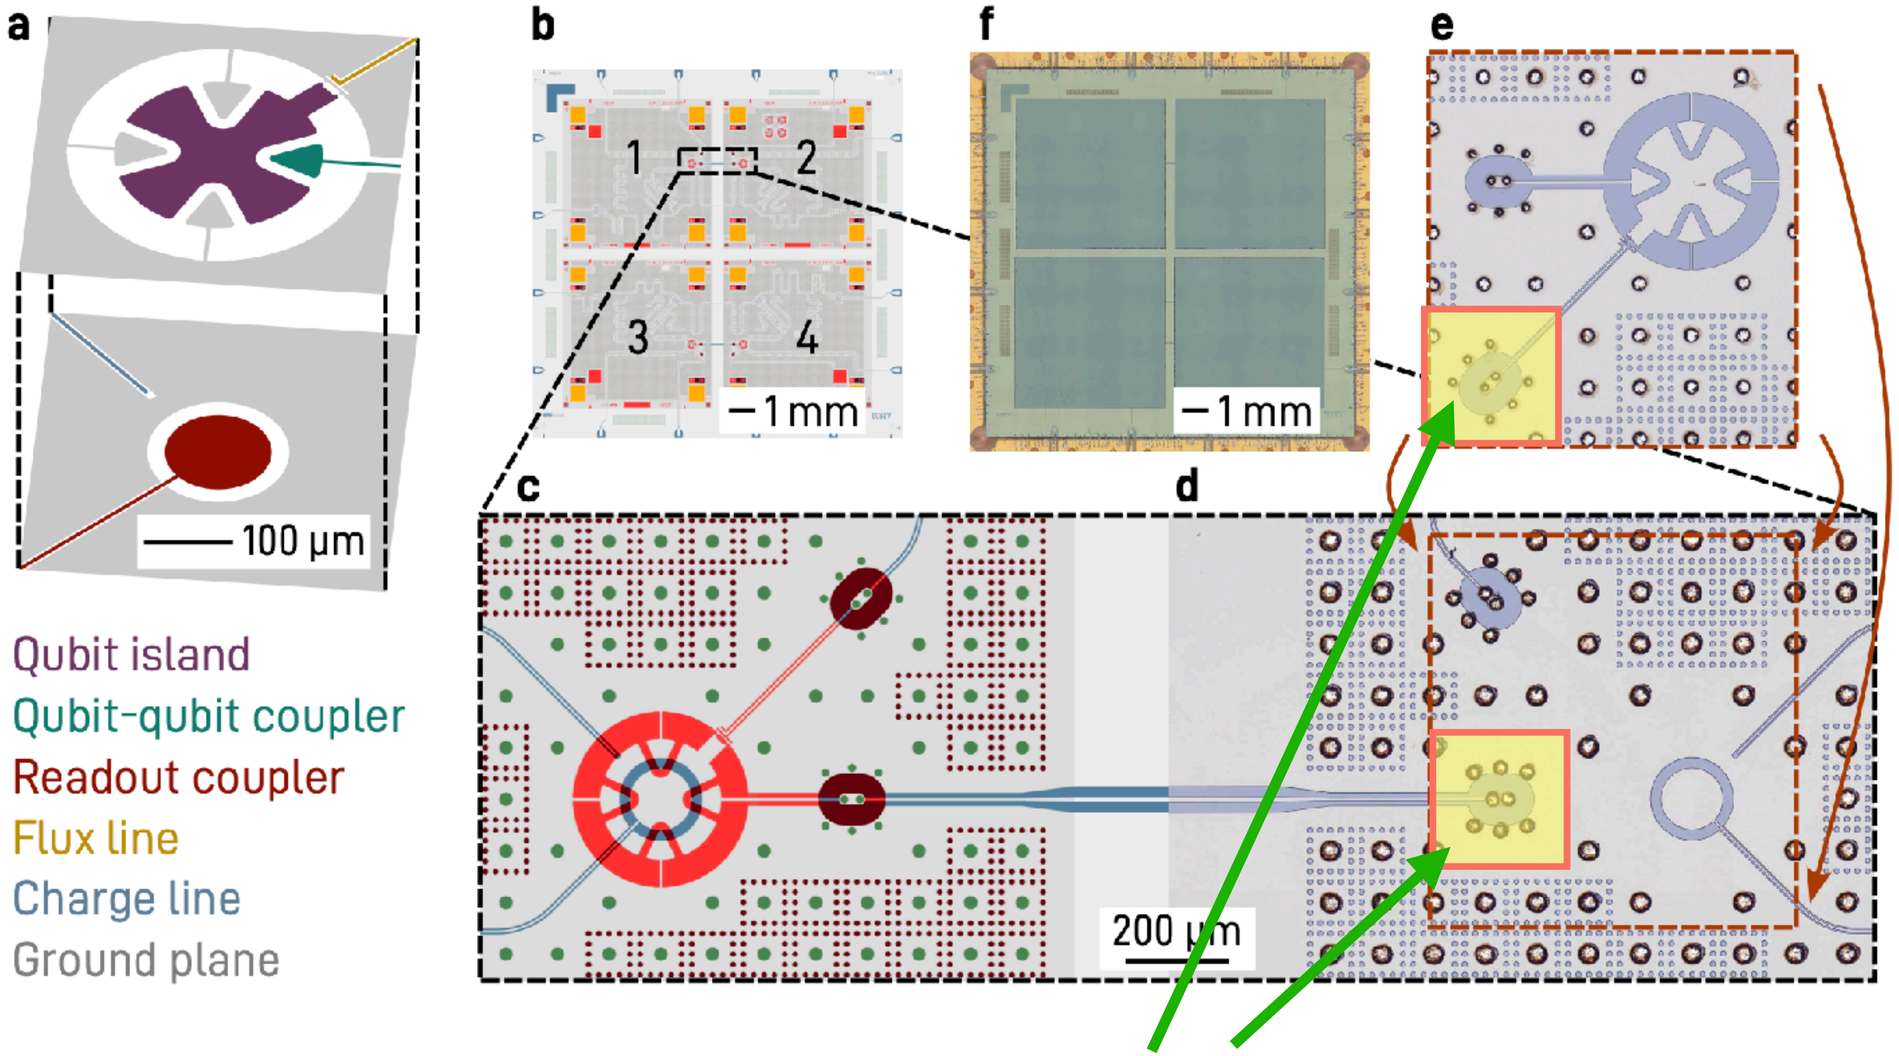

G. J. Norris et al., arXiv:2503.12603

Quantum communication

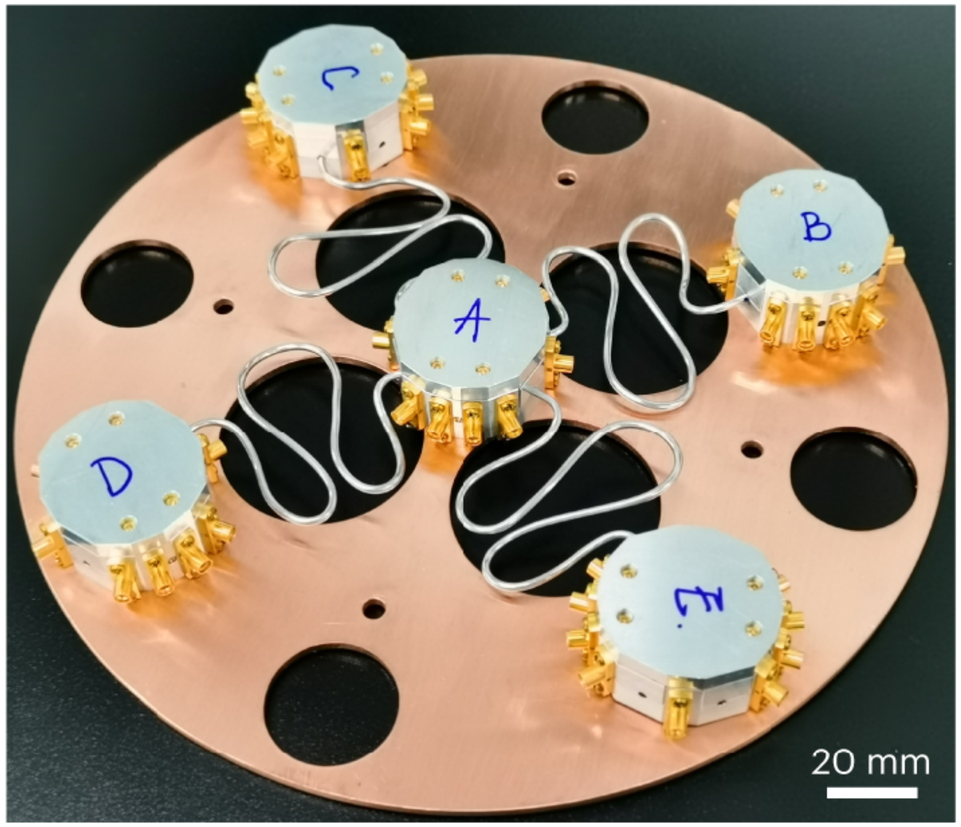

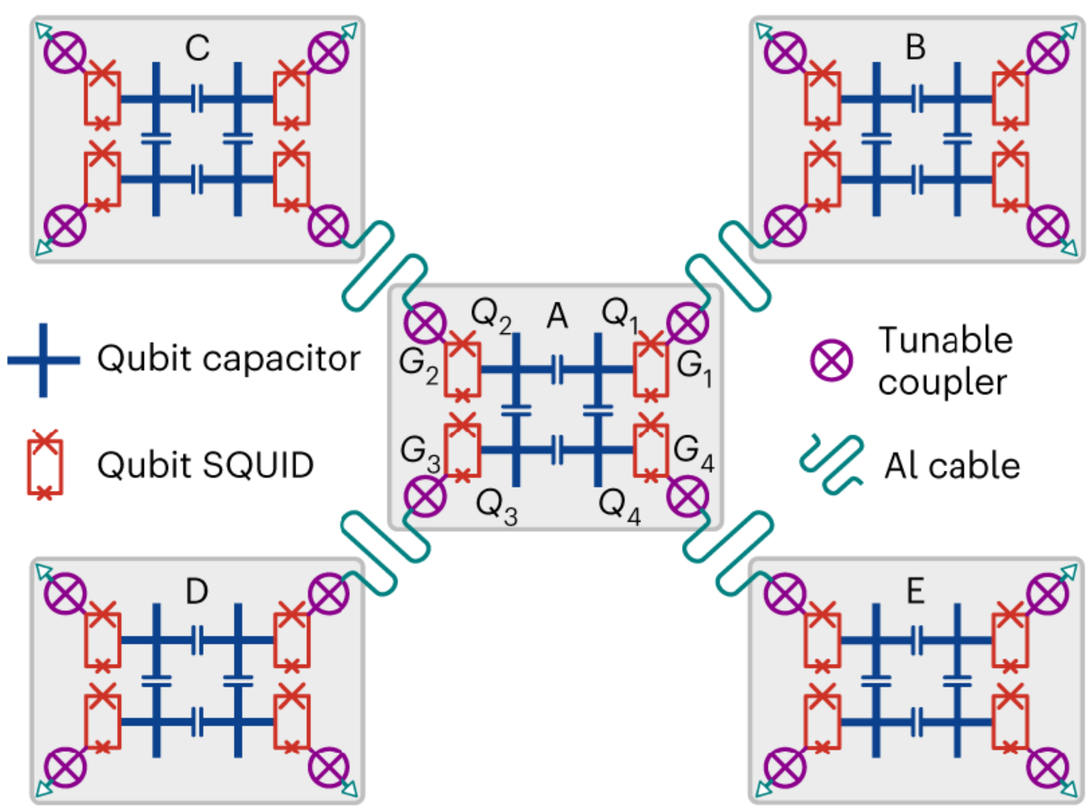

J. Niu *et al.*, Nat. Electron. **6**, 235 (2023)
Y. Jhong *et al.*, Nature **590**, 571 (2021)

Classical communication

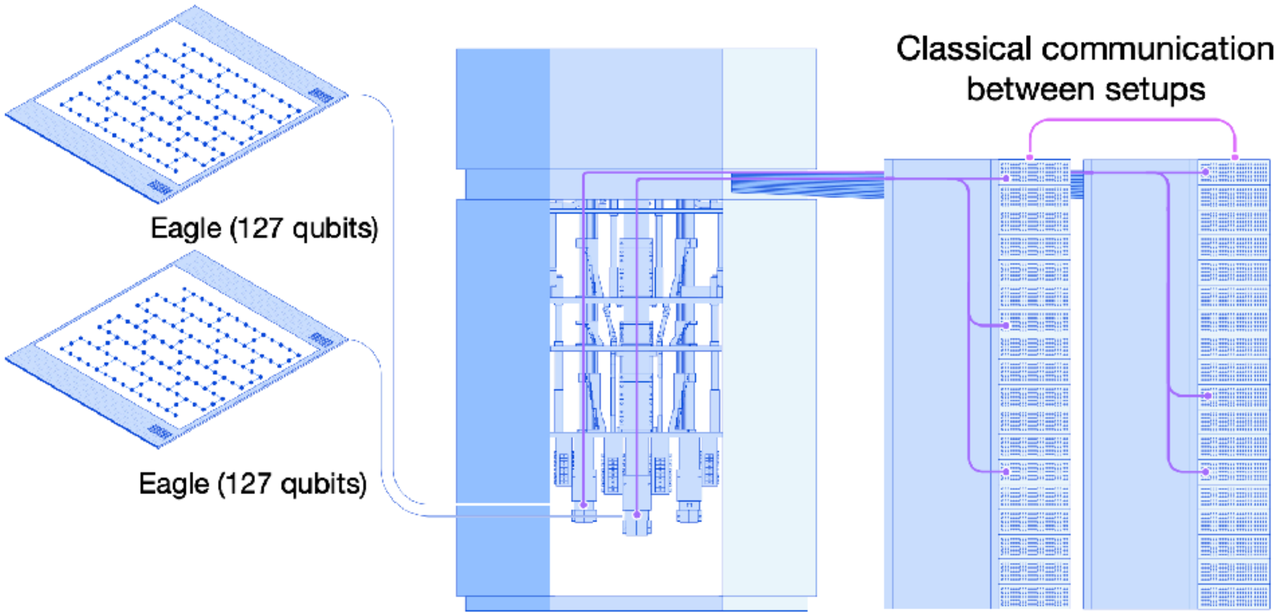

Circuit knitting

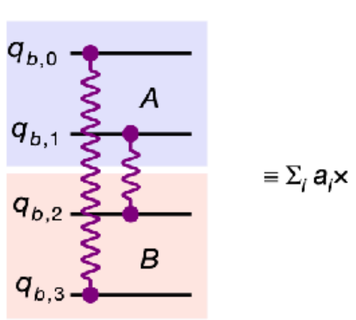

$\equiv \Sigma_i a_i \times$

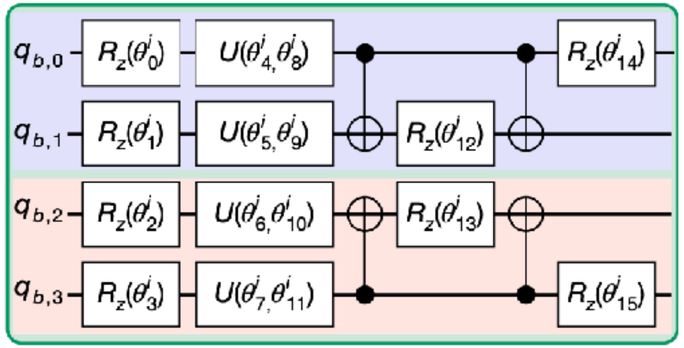

$$CNOT = (I \otimes I + Z \otimes I + I \otimes Z - Z \otimes X)/2$$

A. C. Vazquez et al., Nature 636, 75 (2024)

Distributed quantum algorithms

Distributed quantum phase estimation algorithm

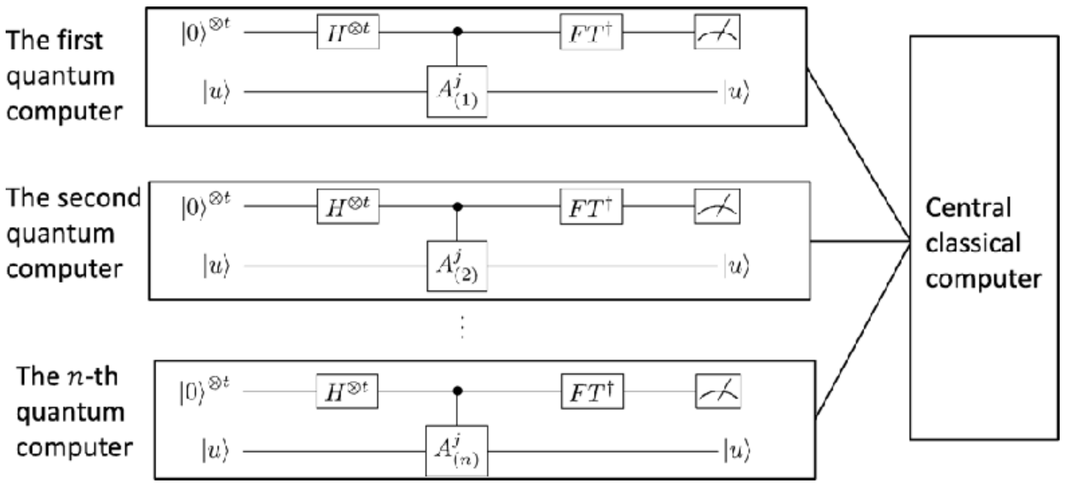

Distributed Shor's factoring algorithm

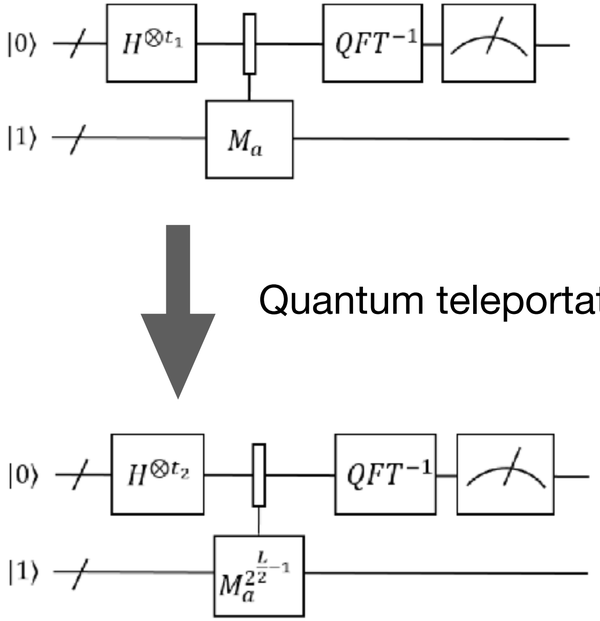

K. Li et al. Inf. Process. Lett. 120, 23 (2017)

L Xiao et al., arXiv:2207.05976v1

Variational quantum algorithms

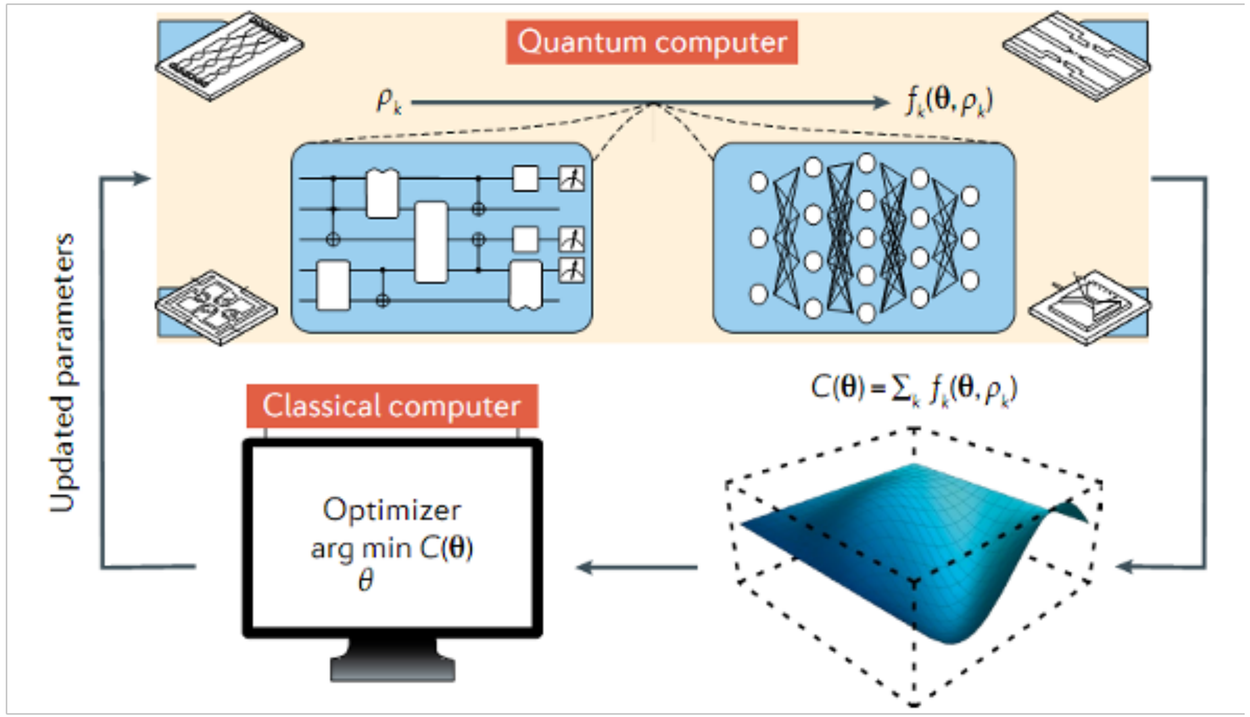

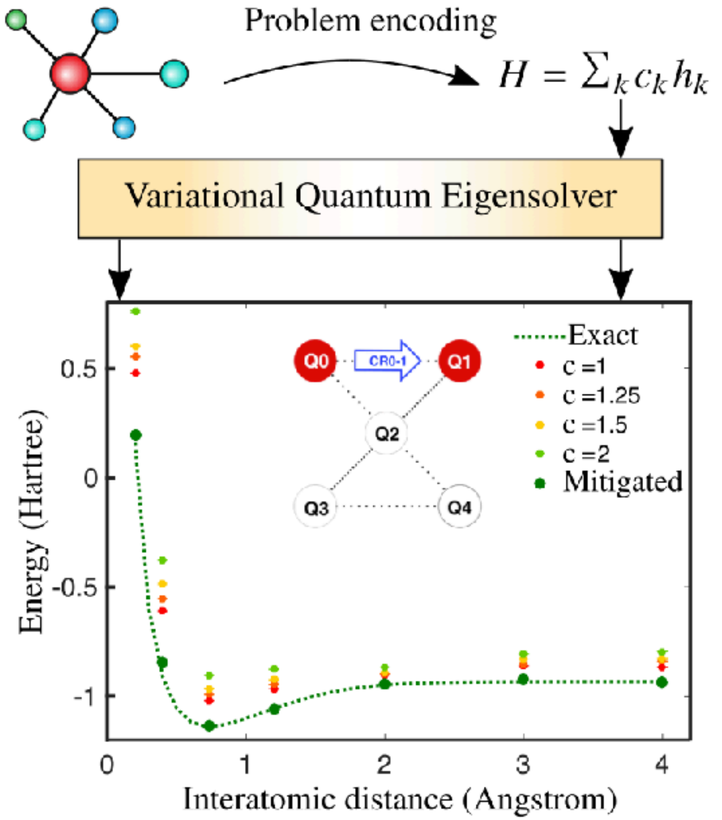

Do not require full circuit universality.

Sufficient expressibility to explore relevant solution spaces is enough!

M. Cerezo et al., *Nat. Rev. Phys.* 3, 625 (2021)

Quantum machine learning

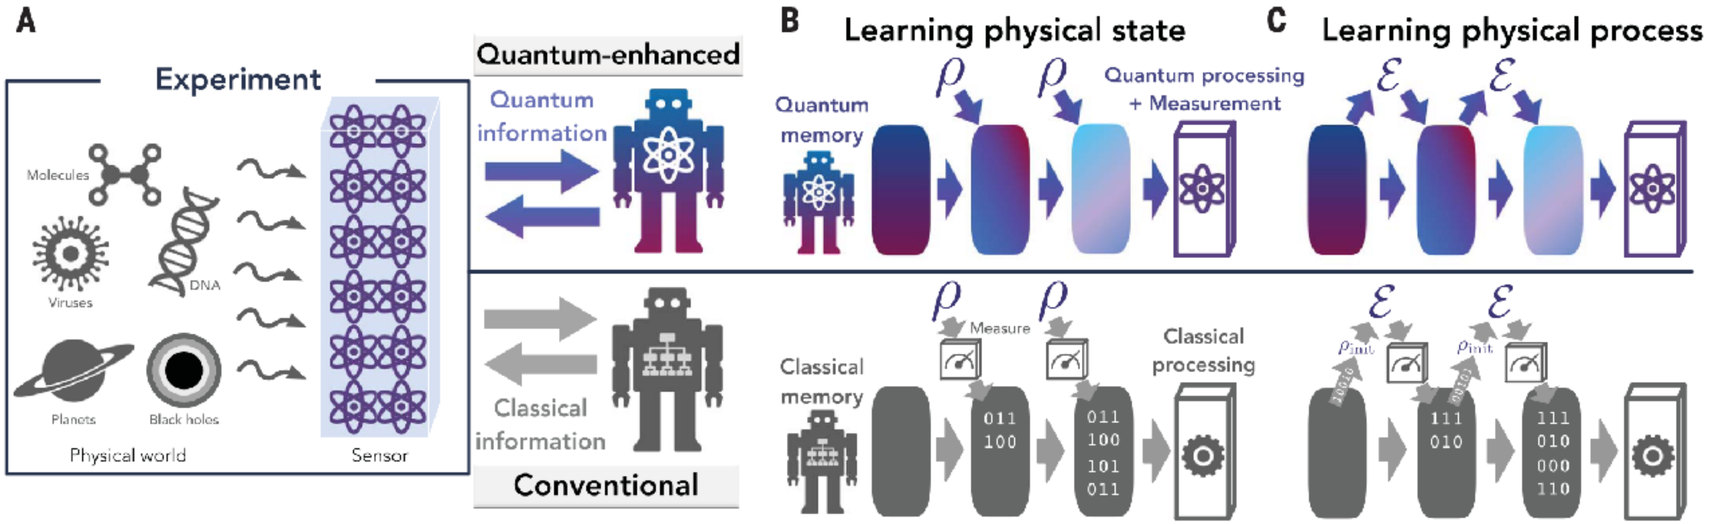

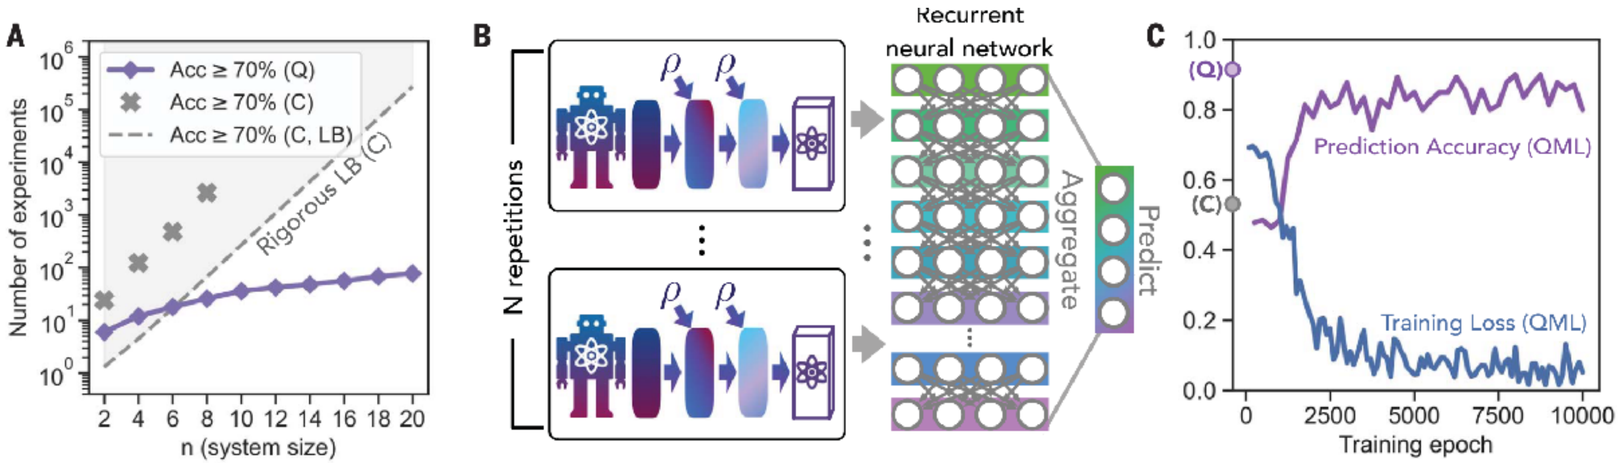

Science 376, 1182–1186 (2022)

Quantum machine learning

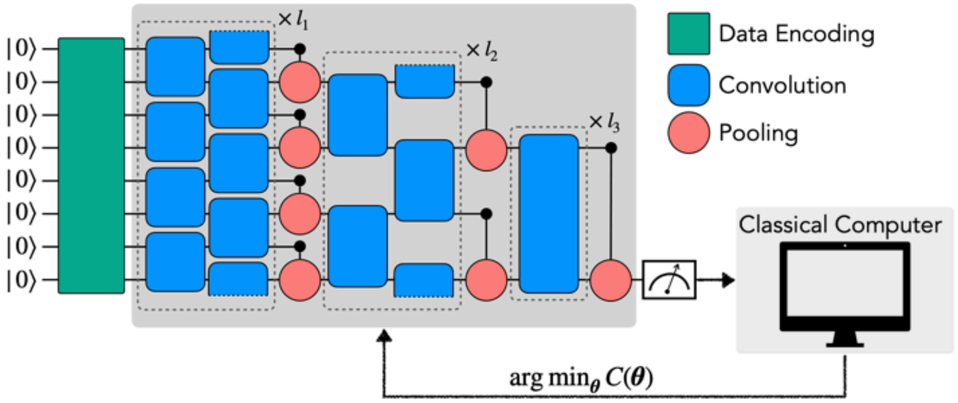

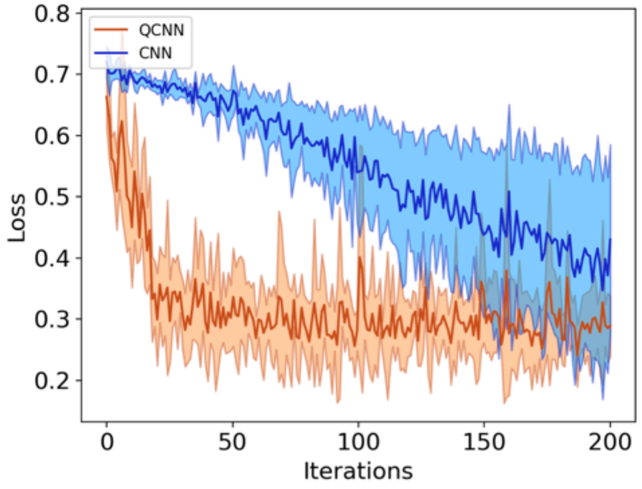

(a) 8-input CNN with AutoEnc vs QCNN

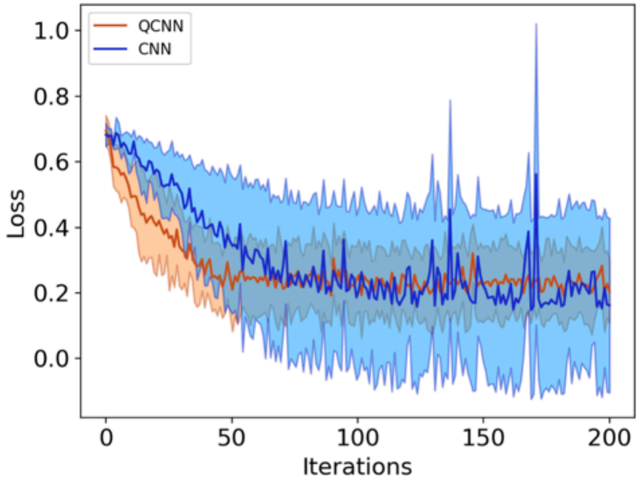

(b) 16-input CNN with PCA vs QCNN

T. Hur, L. Kim, D. K Park,
Quantum Machine Intelligence **4**, 3 (2022)

Distributed quantum machine learning

Binary classification

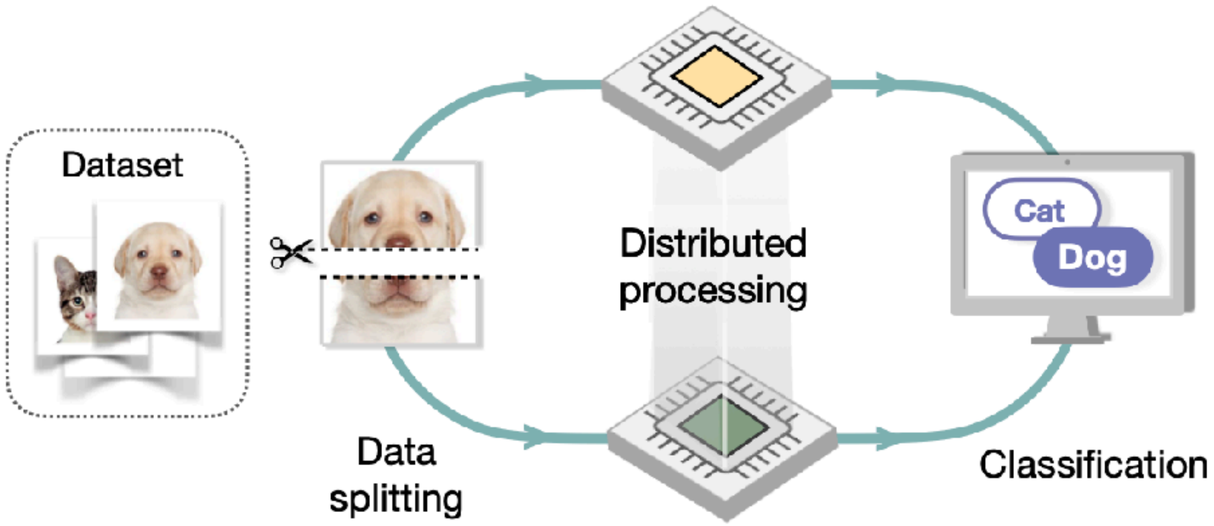

Assessing the efficacy of resources for distributed quantum computing

Communication resources

1. *Classical bits*
2. *Ebits (Bell sources)*
3. *Qubits*

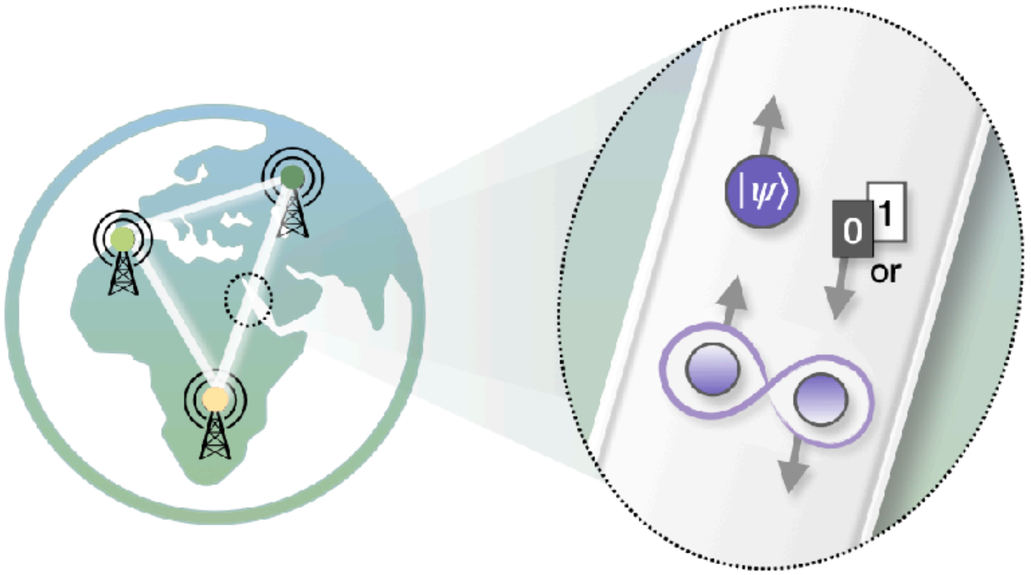

Classical communication

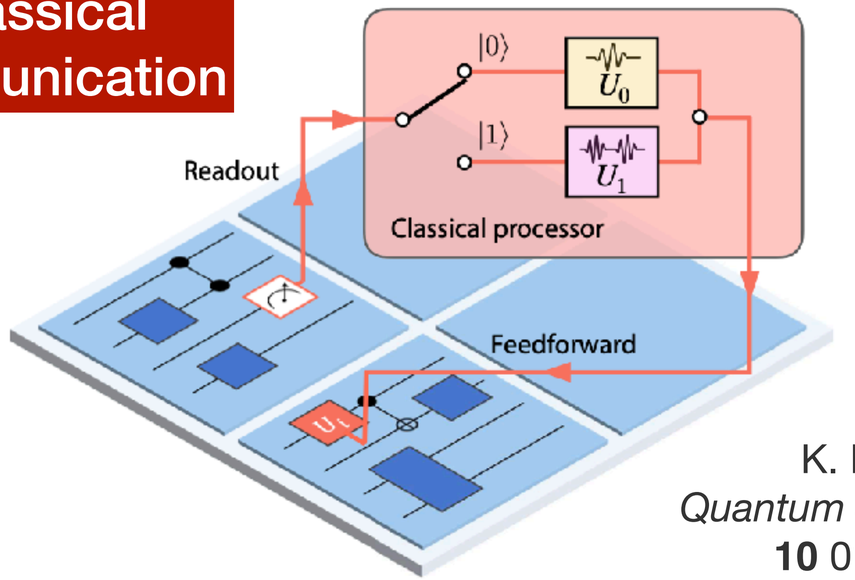

K. Hwang *et al.*,
Quantum Sci. Technol.
10 015059 (2025)

Entanglement

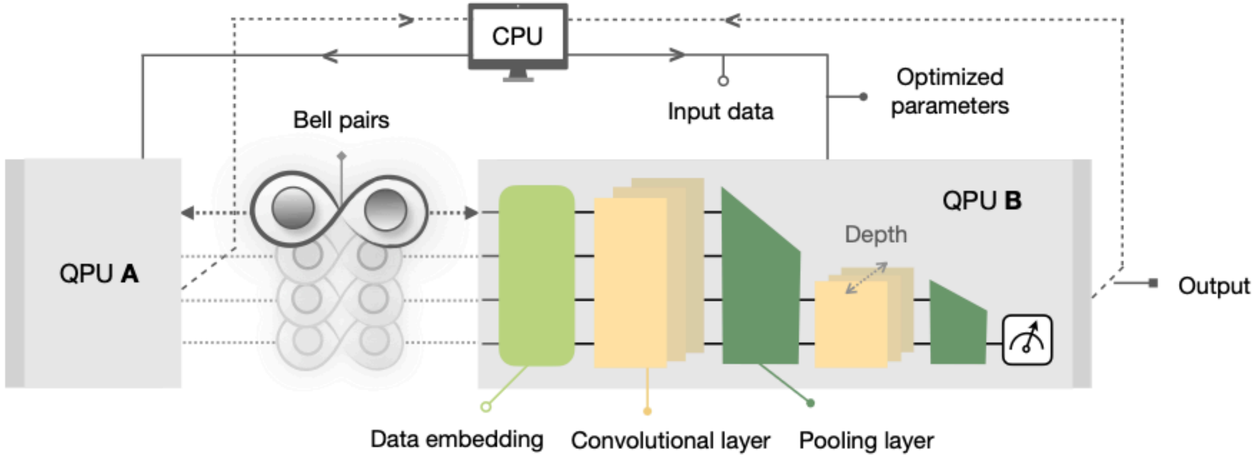

Y. Kim *et al.*, in preparation

Distributed quantum machine learning via classical communication

K. Hwang[1,2,3], H.-T. Lim[1], Y.-S. Kim[1], D. K. Park[2], and **Y. Kim**[3]

Quantum Sci. Technol. **10**, 015059 (2025)

[1] KIST, [2] Yonsei University, [3] Korea University

Quantum machine learning — Binary classification

Data embedding

Classical data => Quantum states

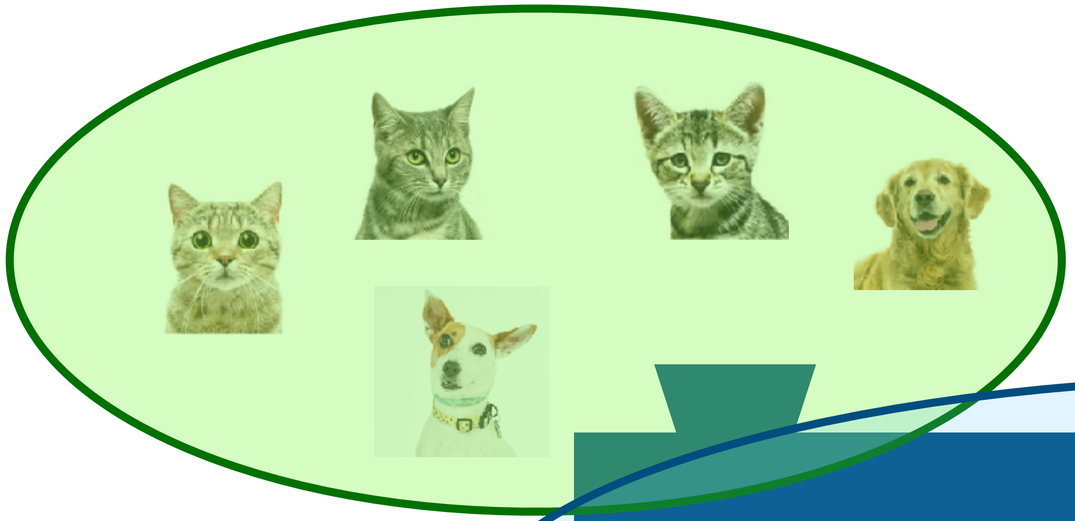

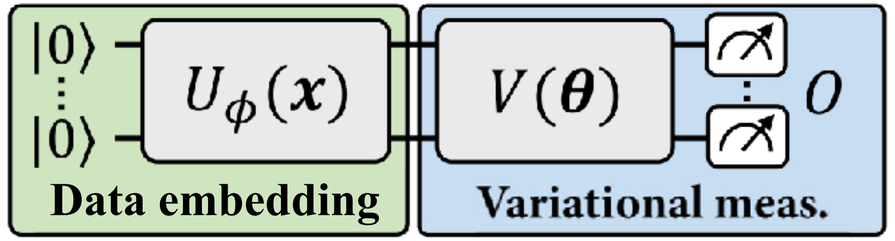

Training network

Optimizing parameterized quantum circuits

Optimizing θ

$f(x\,|\,\theta)$

"Cat" or "Dog"

Classification

Expectation value of an observable

Quantum convolutional neural networks

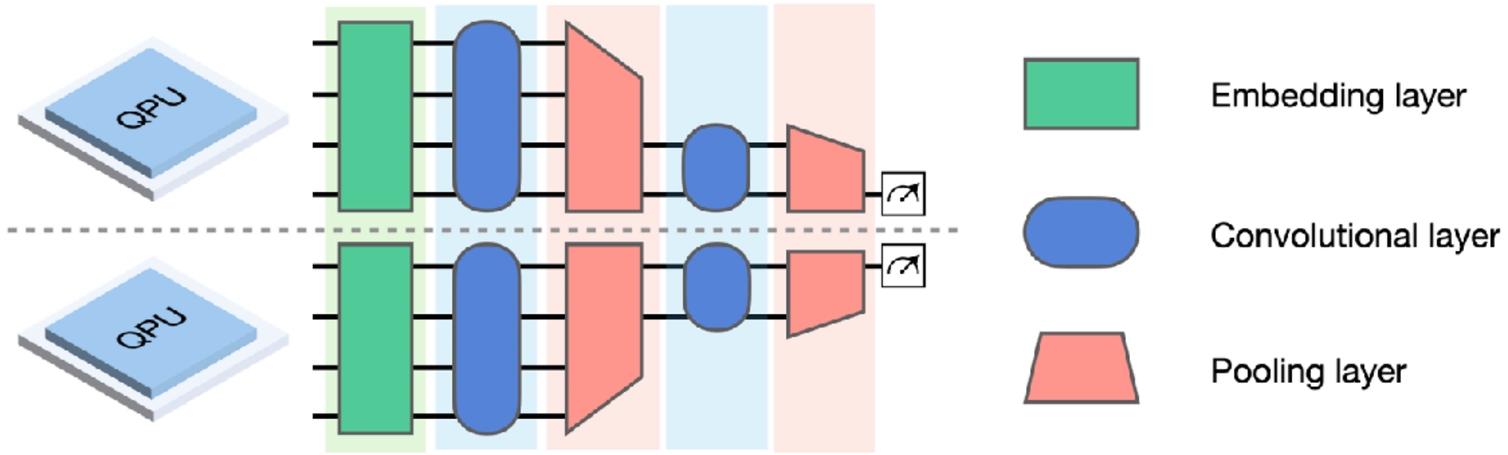

Classical communication implementation

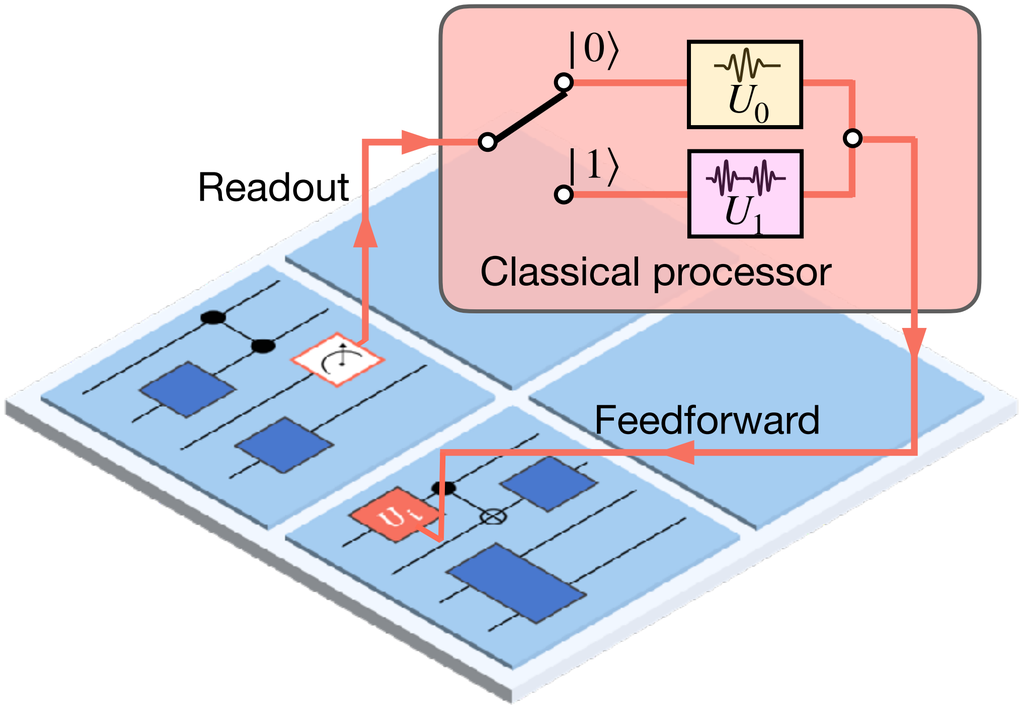

Mid-Circuit Measurements and Feedforward

	Conditional Pulse		Control Flow		Param. Amplitude Update	
	Distributed	Aggregated	Distributed	Aggregated	Distributed	Aggregated
Channel #	Active Reset		RUS Active Reset		π-pulse Calibration	
1	224 ns	224 ns	276 ns	276 ns	272 ns	272 ns
20	238 ns	424 ns	320 ns	460 ns	296 ns	472 ns
50	238 ns	608 ns	320 ns	936 ns	296 ns	720 ns

200 ns << Coherence time

Results

NC: No communication
QC: Quantum communication
CC: Classical communication

Data: 8-bit vectors
Label: Red / Blue

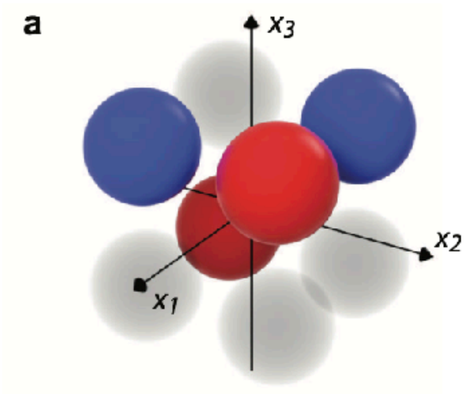

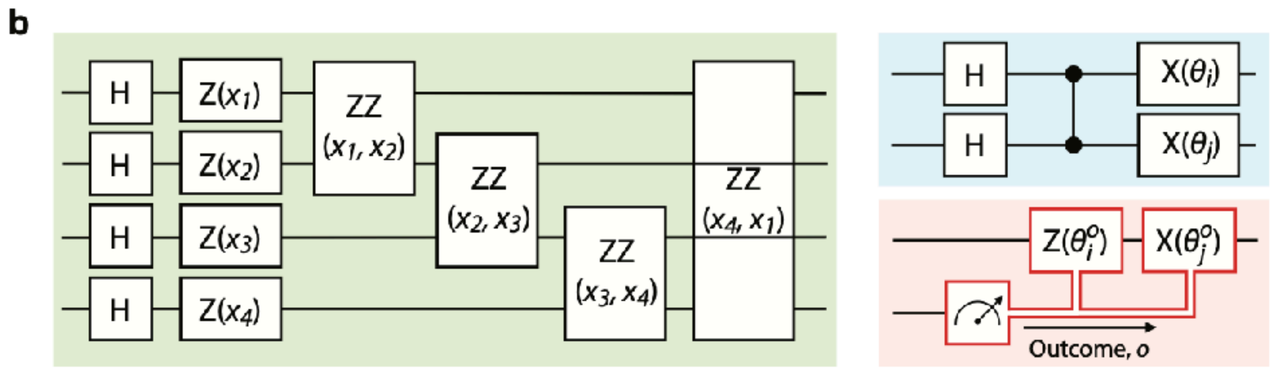

Circuit capacity analysis

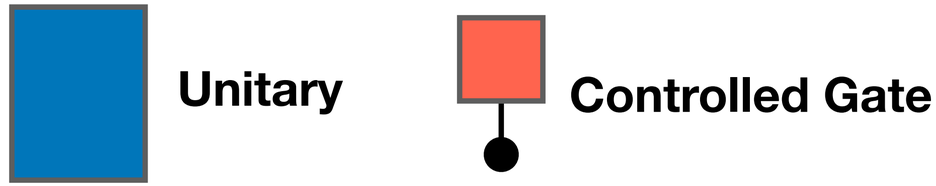

Set of circuits

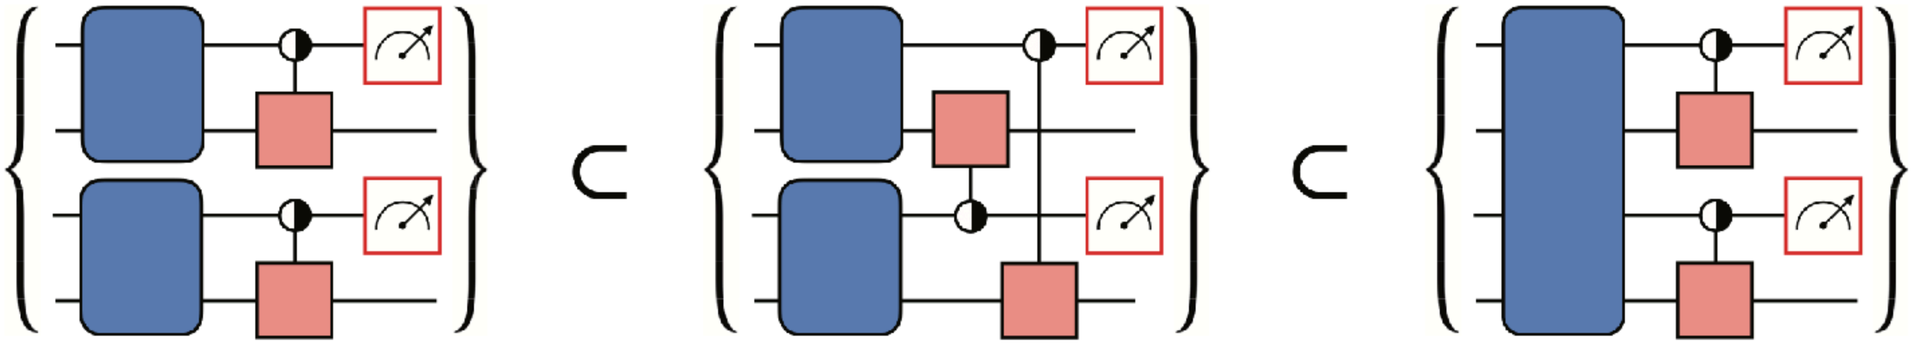

- Control operation is absorbed by preceding unitary operation.
- Control operation between processors cannot be absorbed into convolutional layer.

Circuit capacity analysis

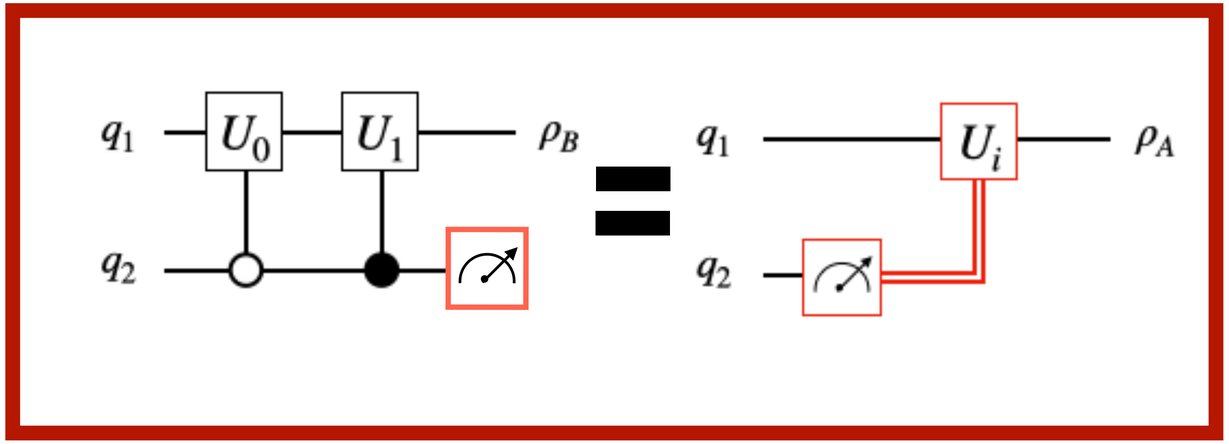

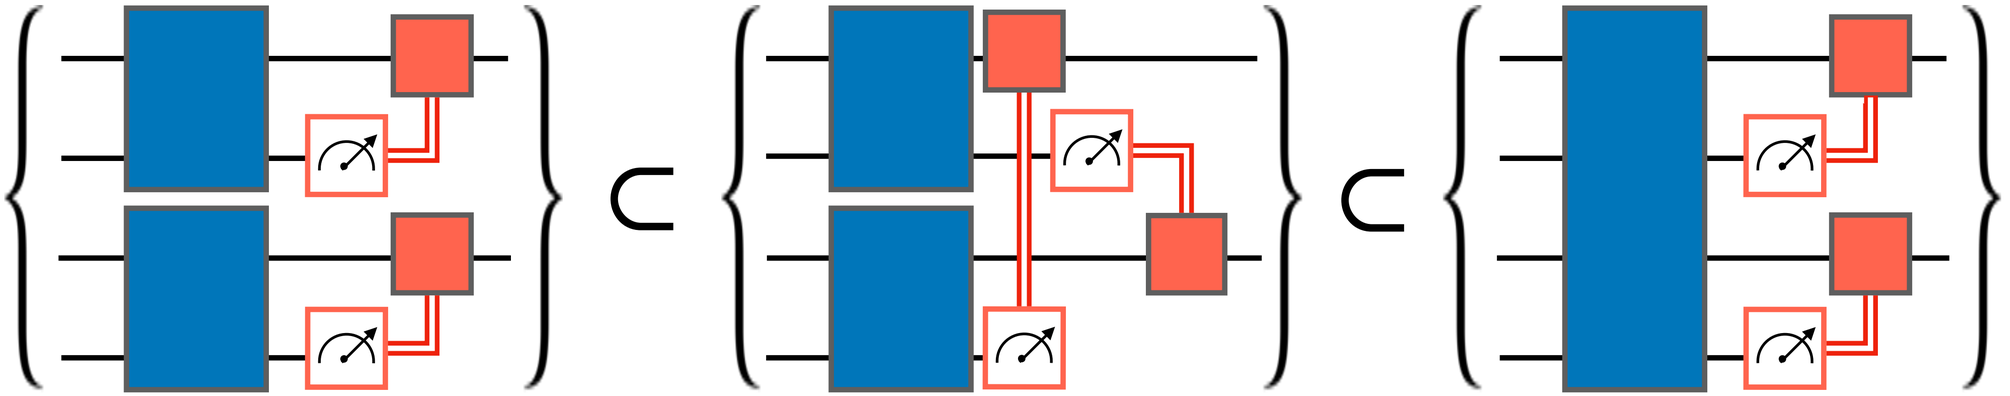

No Communication ⊂ Classical Communication ⊂ Quantum Communication

Effective dimension

# of Effective parameters: 1

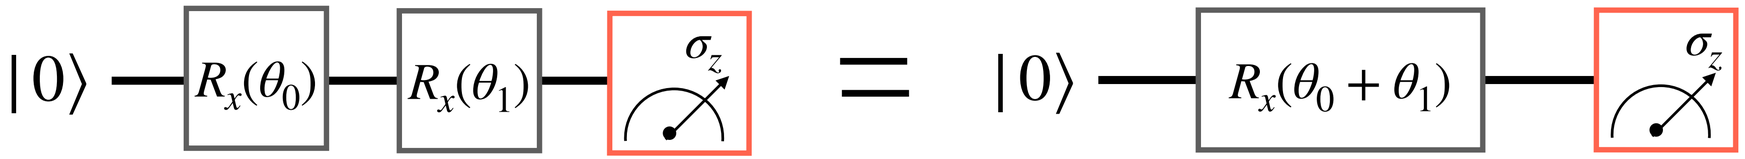

# of Effective parameters: 2

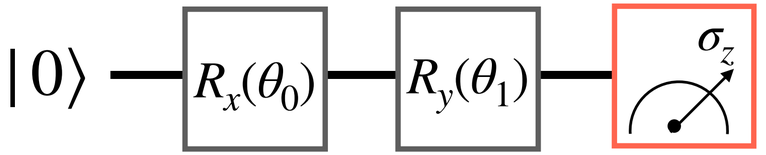

Effective dimension

Fisher information:

$$F(\theta) = \mathbb{E}_{(x,y)\sim p}\left[\frac{\partial}{\partial\theta}\log p(x,y;\theta)\frac{\partial}{\partial\theta}\log p(x,y;\theta)^{\top}\right] = \begin{pmatrix} F_{00} & F_{01} & \cdots \\ F_{10} & F_{11} & \\ \vdots & & \ddots \end{pmatrix}$$

Effective dimension:

$$d_{r,n} = 2\frac{log(\frac{1}{V_\theta}\int_\theta \sqrt{det(I + \frac{rn}{2\pi logn}\hat{F})}d\theta}{log(\frac{rn}{2\pi logn})}$$

$$\lim_{n\to\infty} d_{r,n} = \max_{\theta}[rank[F]]$$

Effective dimension

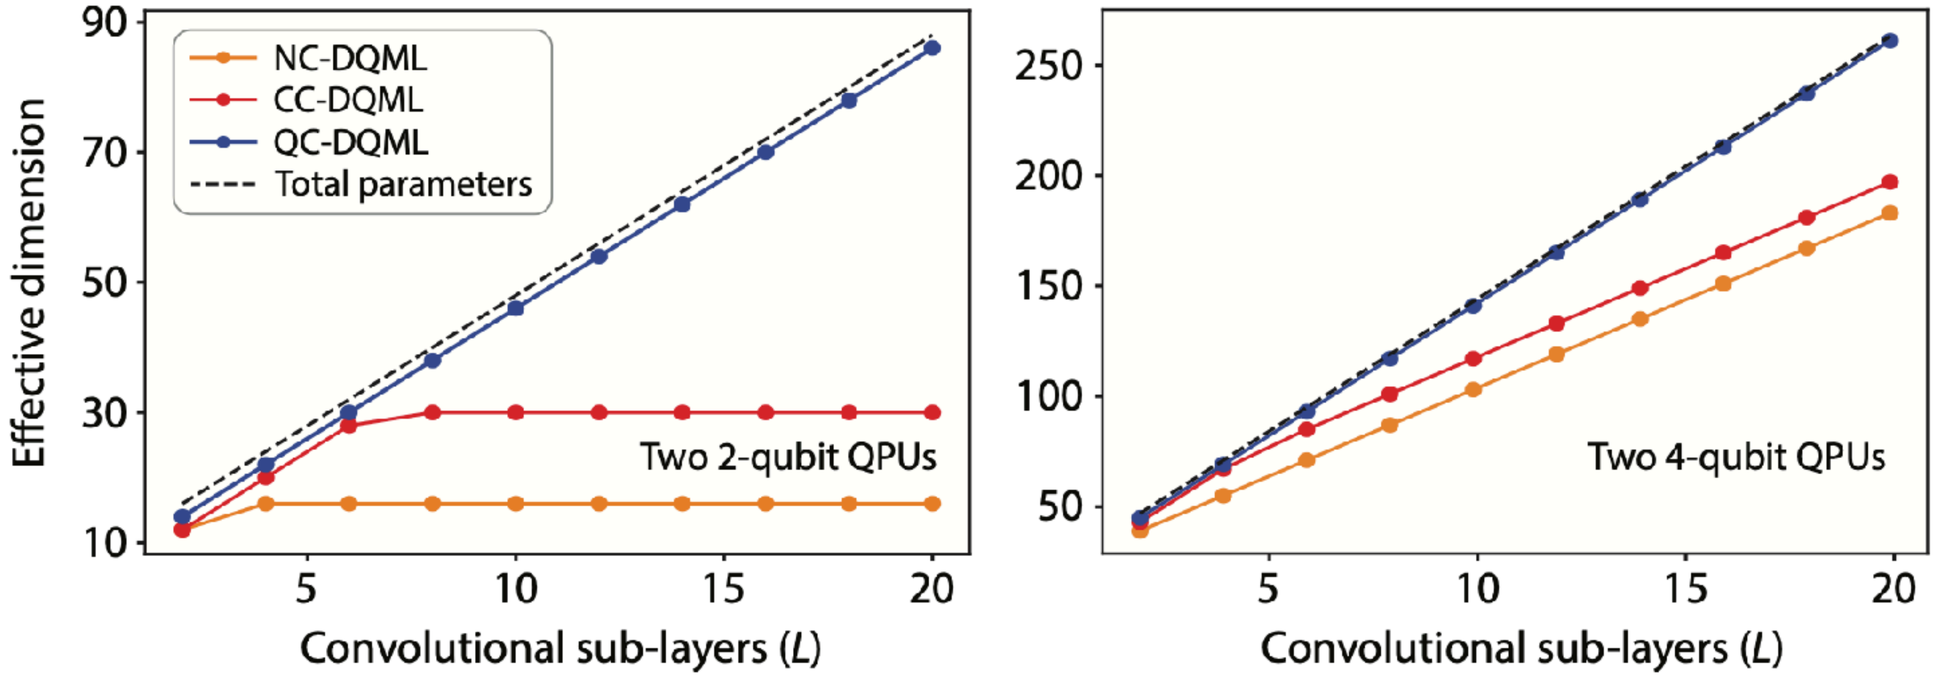

In the shallow-depth regime,
classical communication can substitute for quantum communication.

Results

NC: No communication
QC: Quantum communication
CC: Classical communication

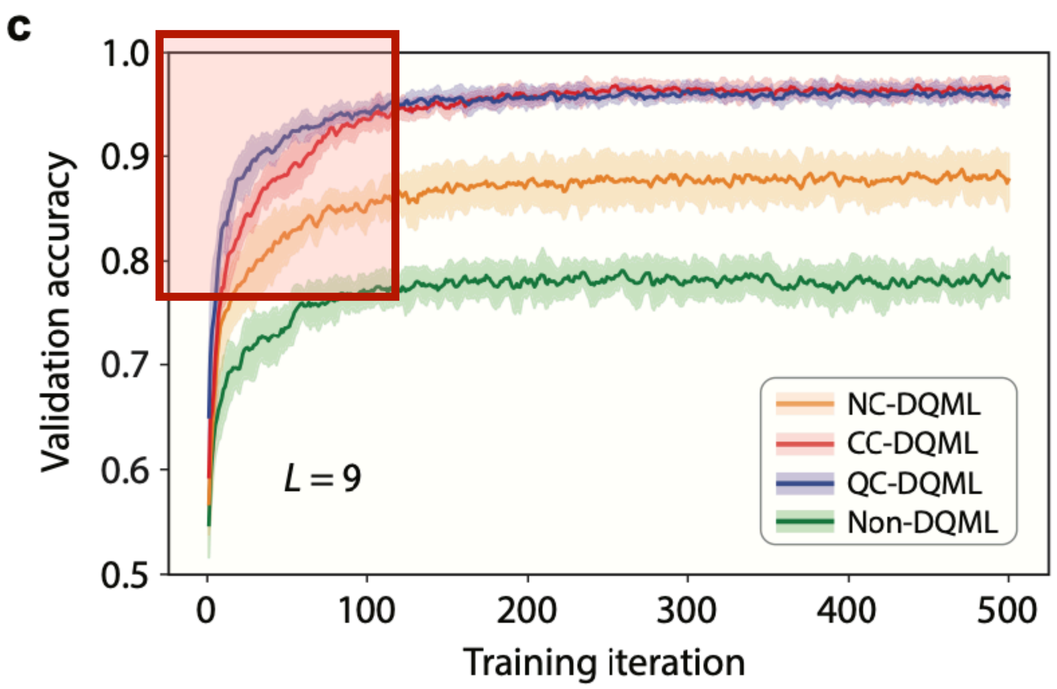

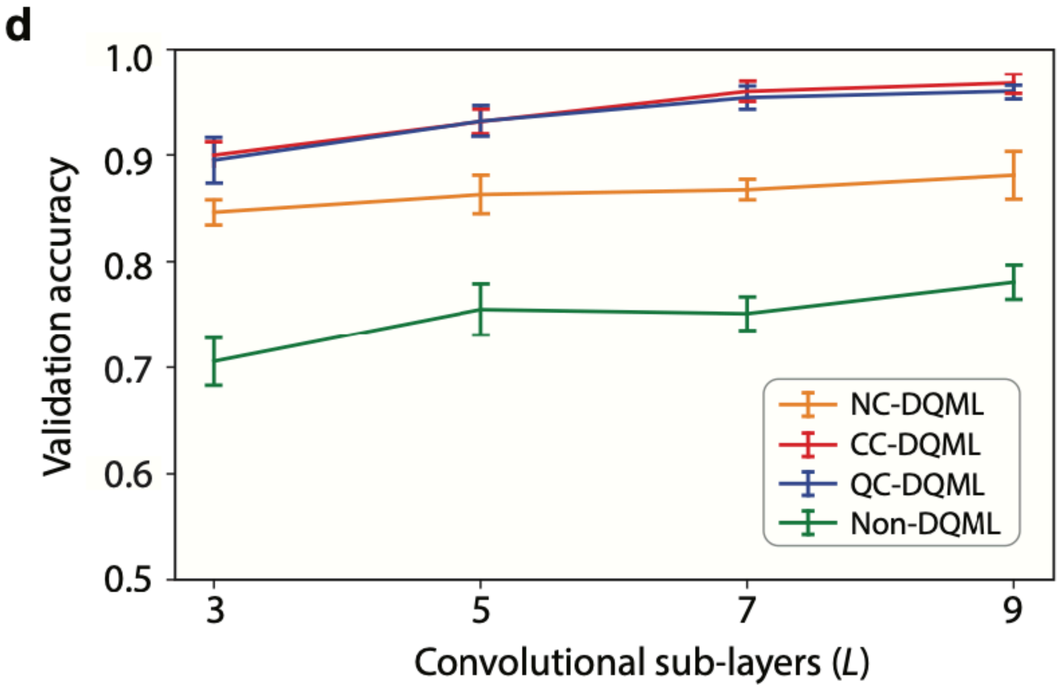

The power of entanglement in distributed quantum machine learning

Y. Kim[1], K. Hwang[1], H. Kwon[2], and **Y. Kim**[1,*]

In preparation

[1] Korea University, [2] KIAS

Quantum internet

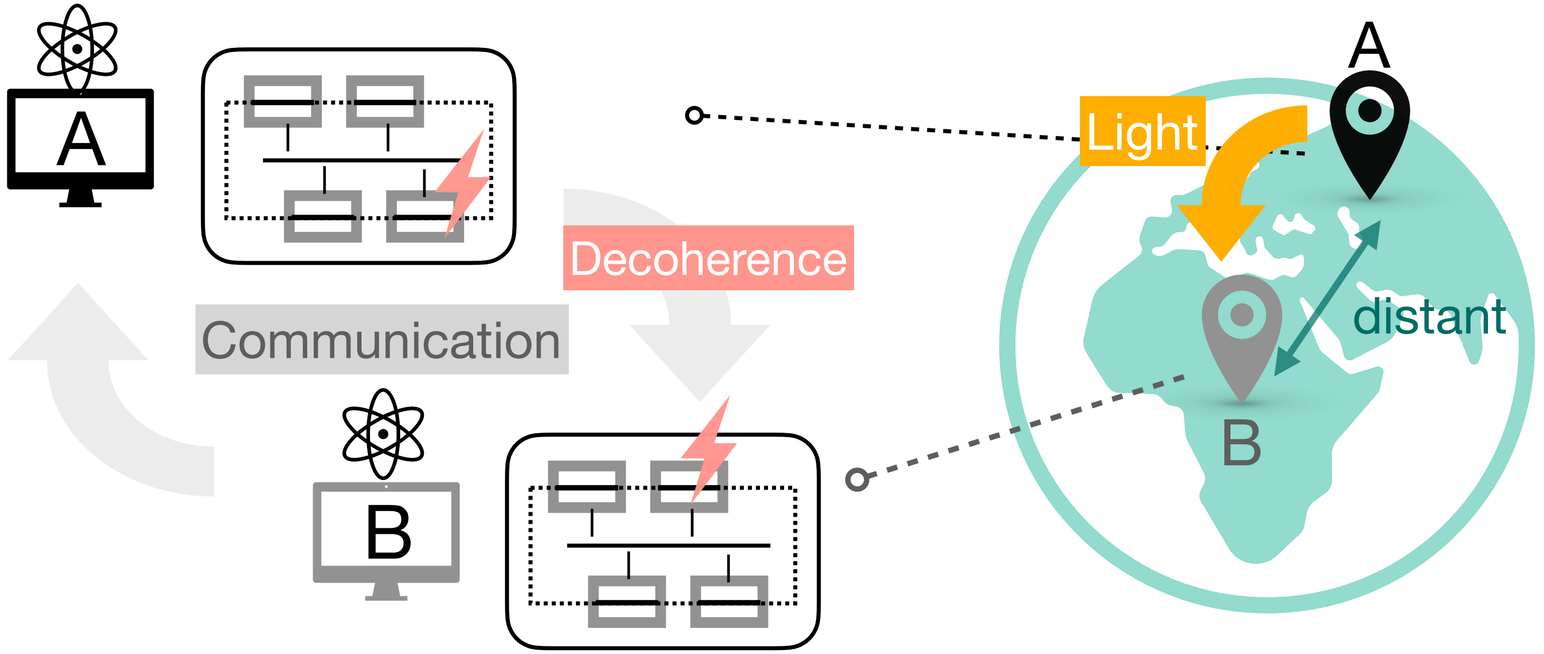

a few hundred kilometers ~ millisecond delays

Communication may not be possible within coherence time

Entanglement distribution

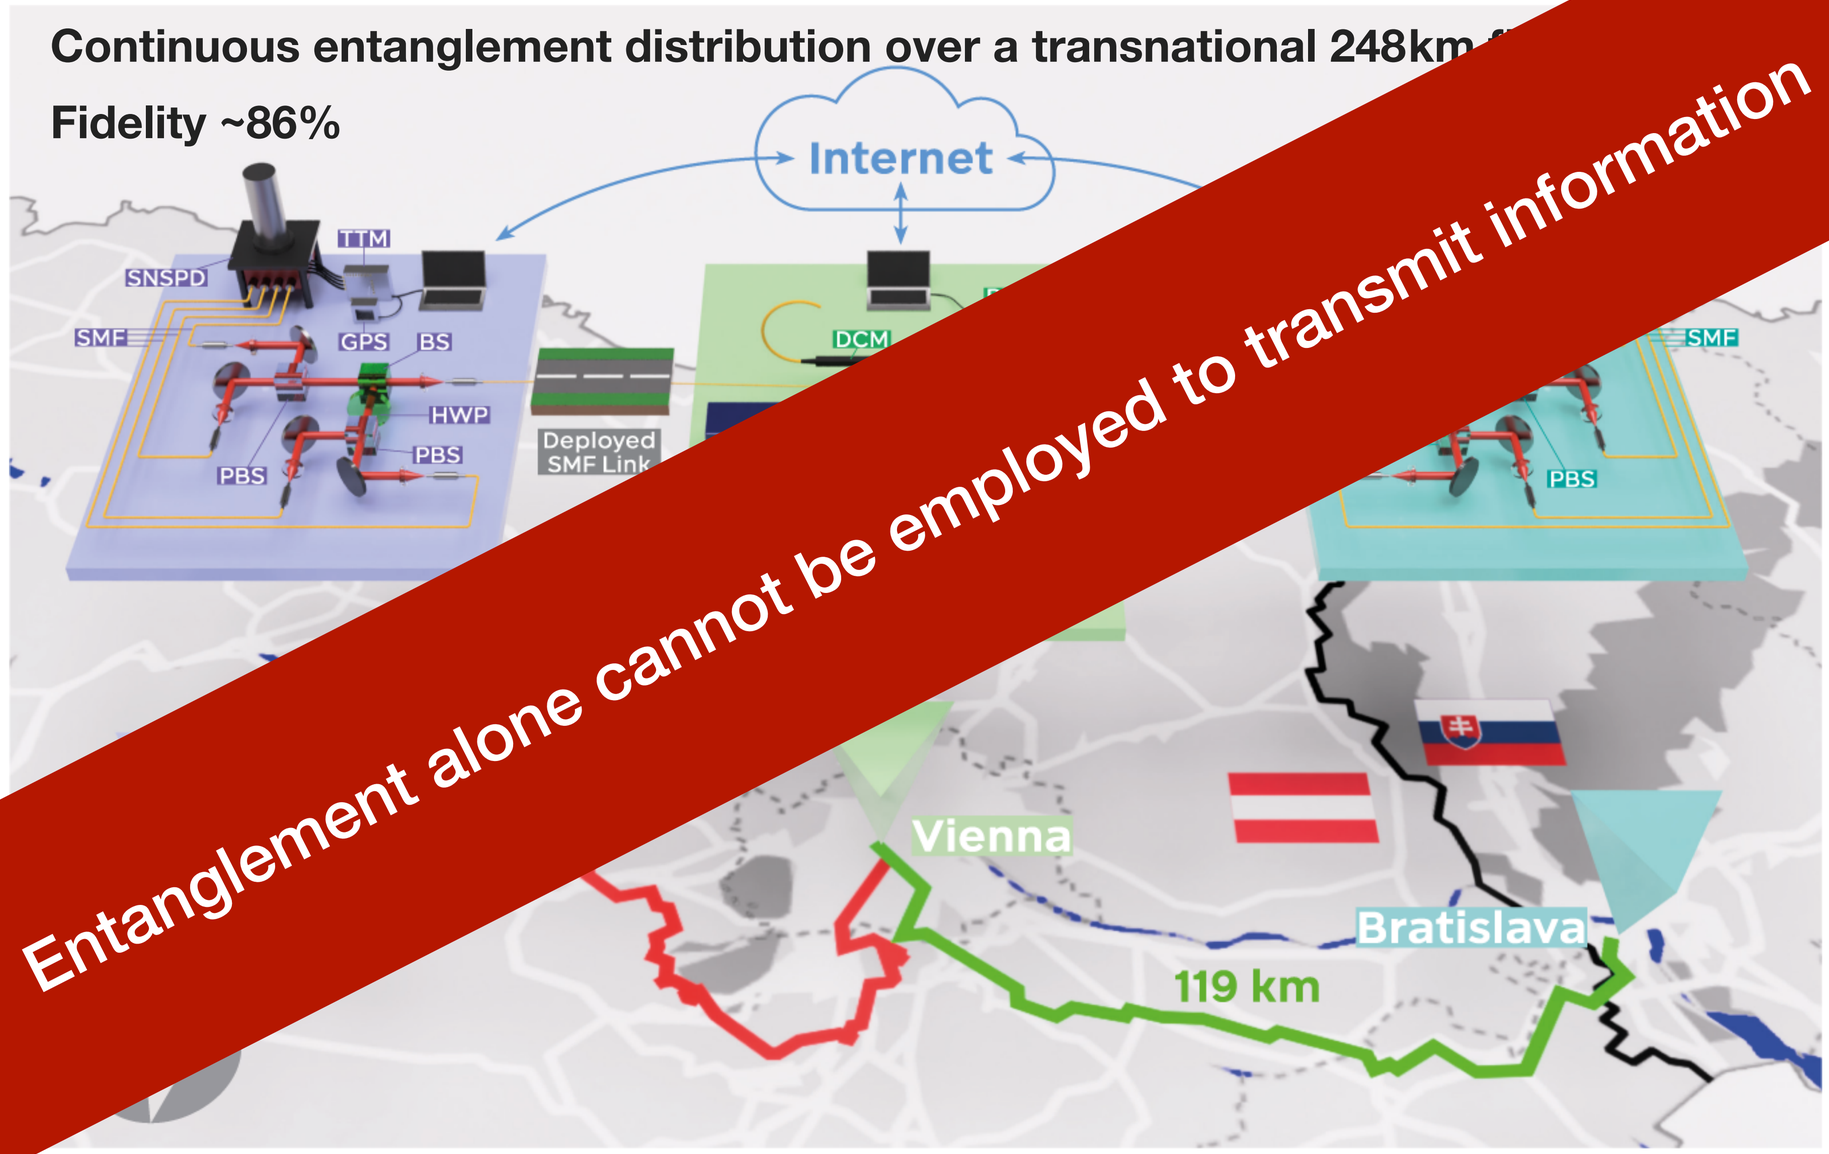

Entanglement alone cannot be employed to transmit information

S. P. Neumann et al., Nat. Commun. 13, 6134 (2022)

Quantum convolutional neural network with Bell sources

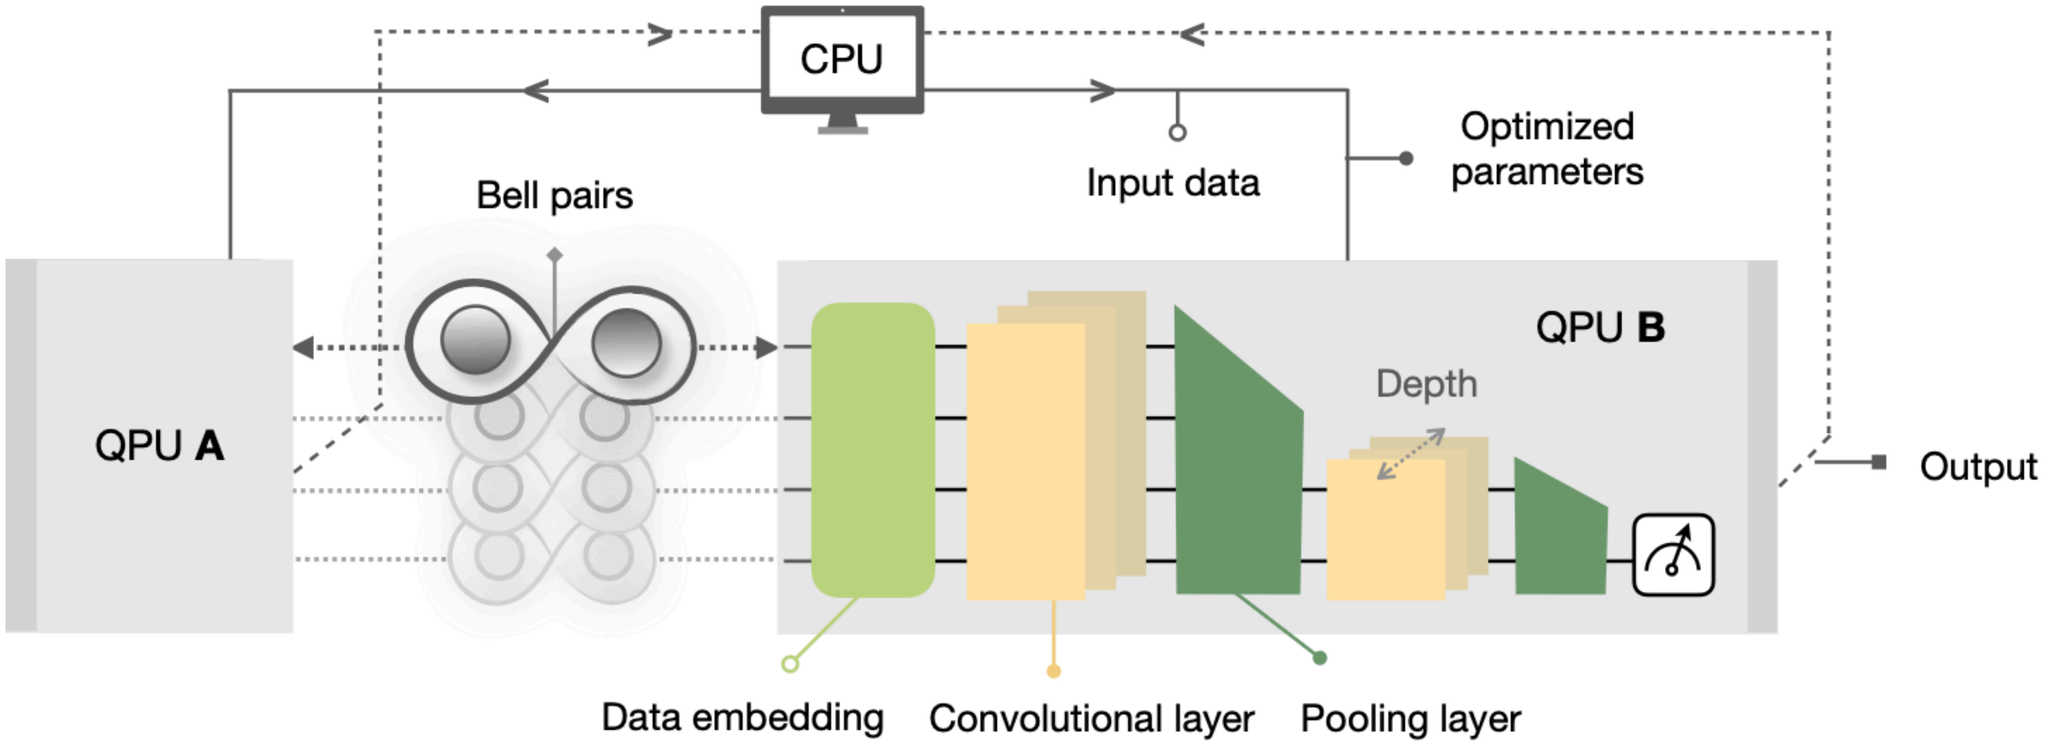

Initial state:

$$|0000\rangle_A|0000\rangle_B \Rightarrow \frac{1}{\sqrt{2}}(|0000\rangle_A|0000\rangle_B + |1000\rangle_A|1000\rangle_B)$$

Results

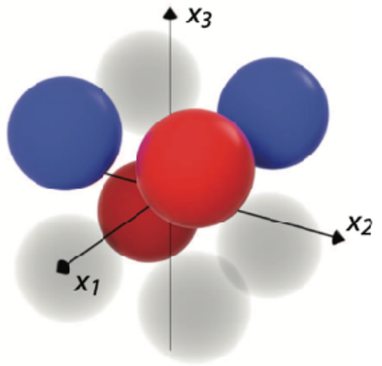

Data: 8-bit vectors
Label: Red / Blue

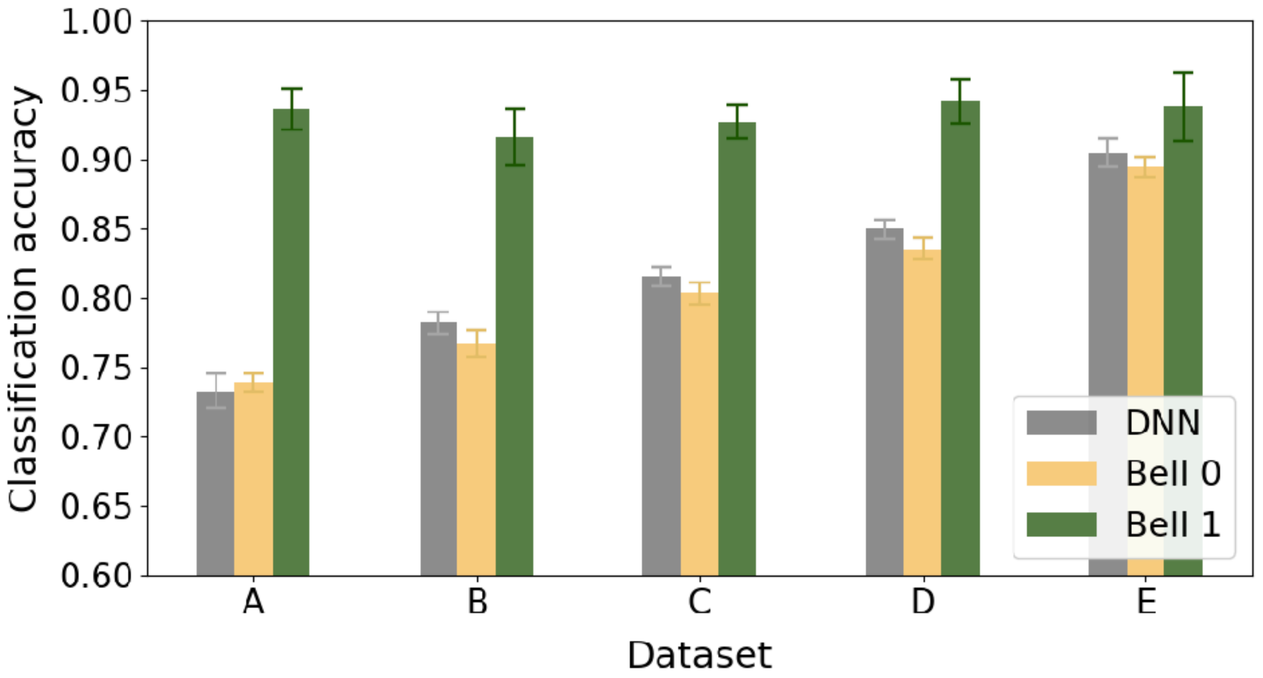

Although performance depends on the dataset,
using a pre-shared Bell source consistently provides an advantage.

Communication complexity

Communication complexity scenario

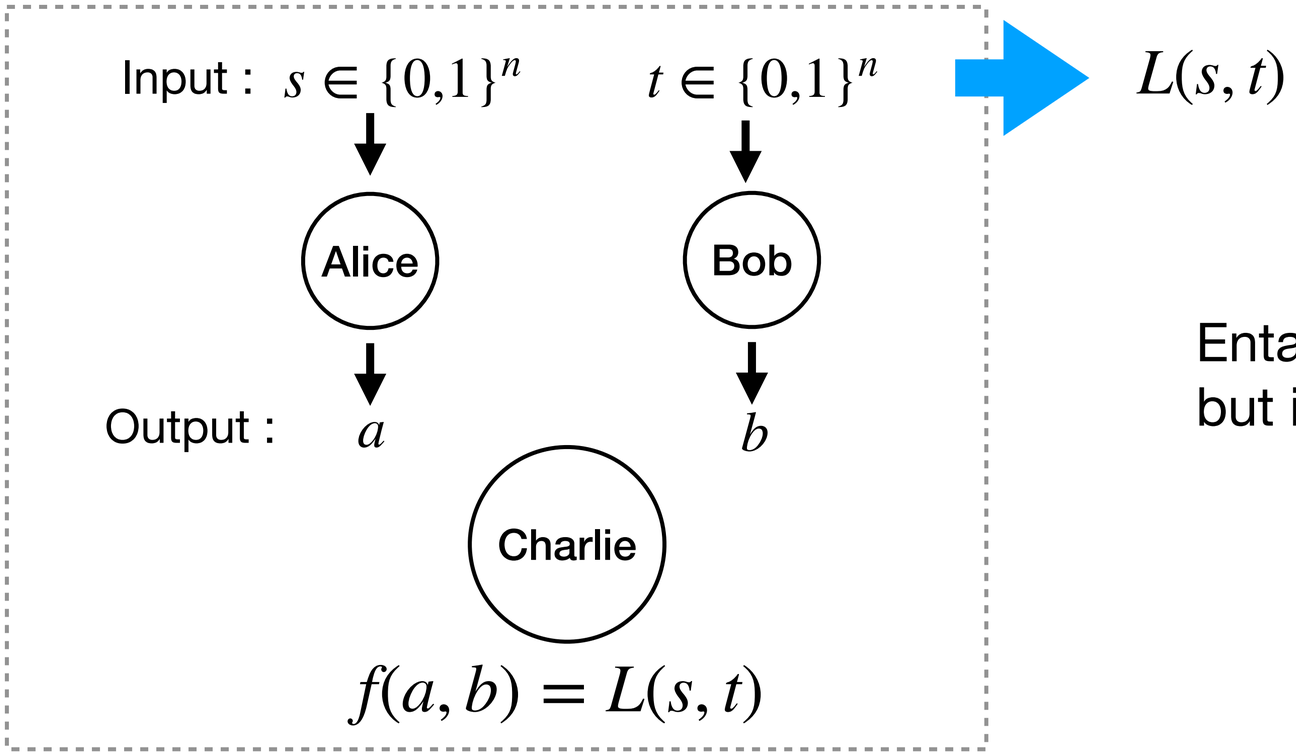

$L(s,t)$

Entanglement cannot carry information,
but it can reduce communication complexity.

Communication complexity - CHSH game

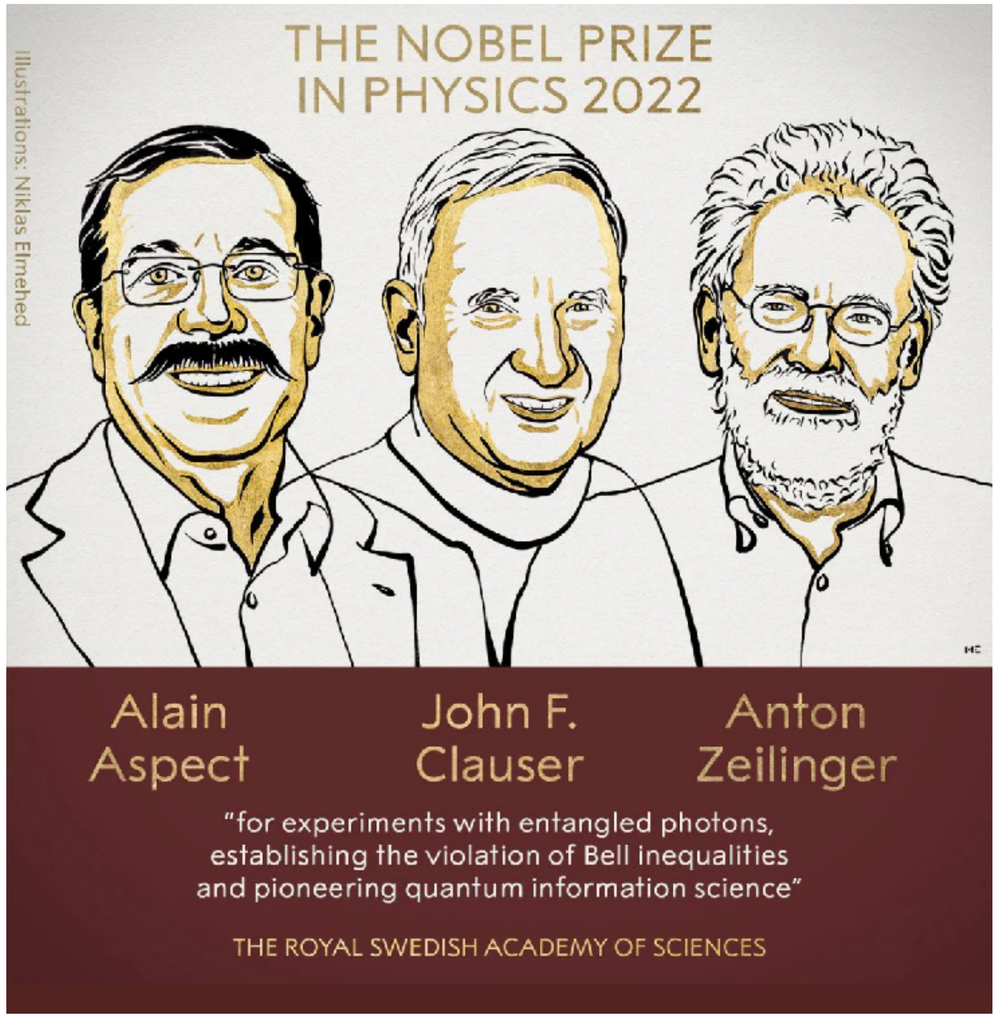

CHSH game

Inputs : s t

Outputs : a b

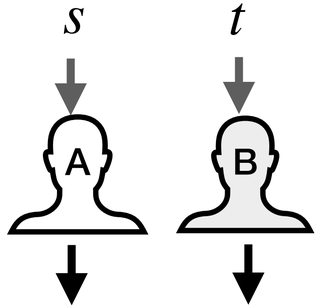

Task : $L = (-1)^{s \wedge t},\ f = ab$

Success rates

Classical correlation: 75%

Quantum correlation: ~85%

Quantum contextuality

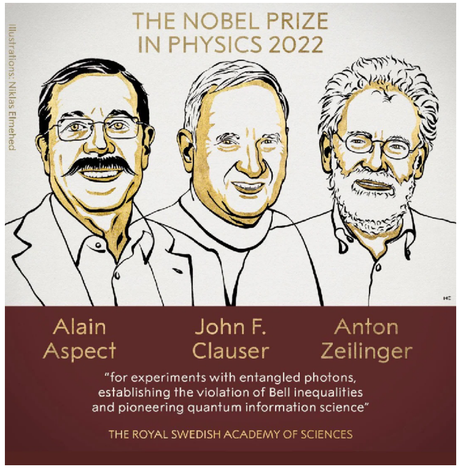

CHSH game

Inputs :

Outputs :

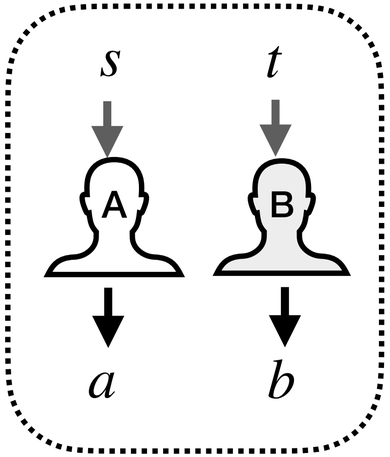

DQML

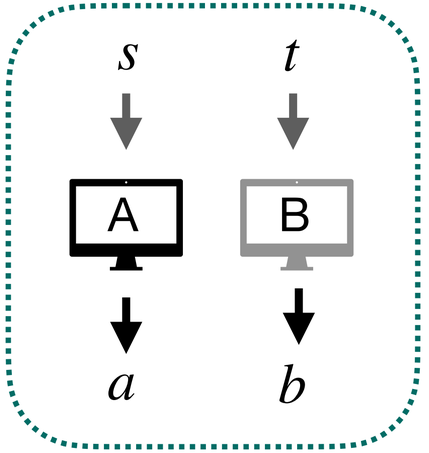

Task :

$$L = (-1)^{s \wedge t}, \quad f = ab$$

$$L(s,t), \quad f(a,b)$$

The advantage of entanglement arises from non-local correlation

Quantum convolutional neural network with Bell sources

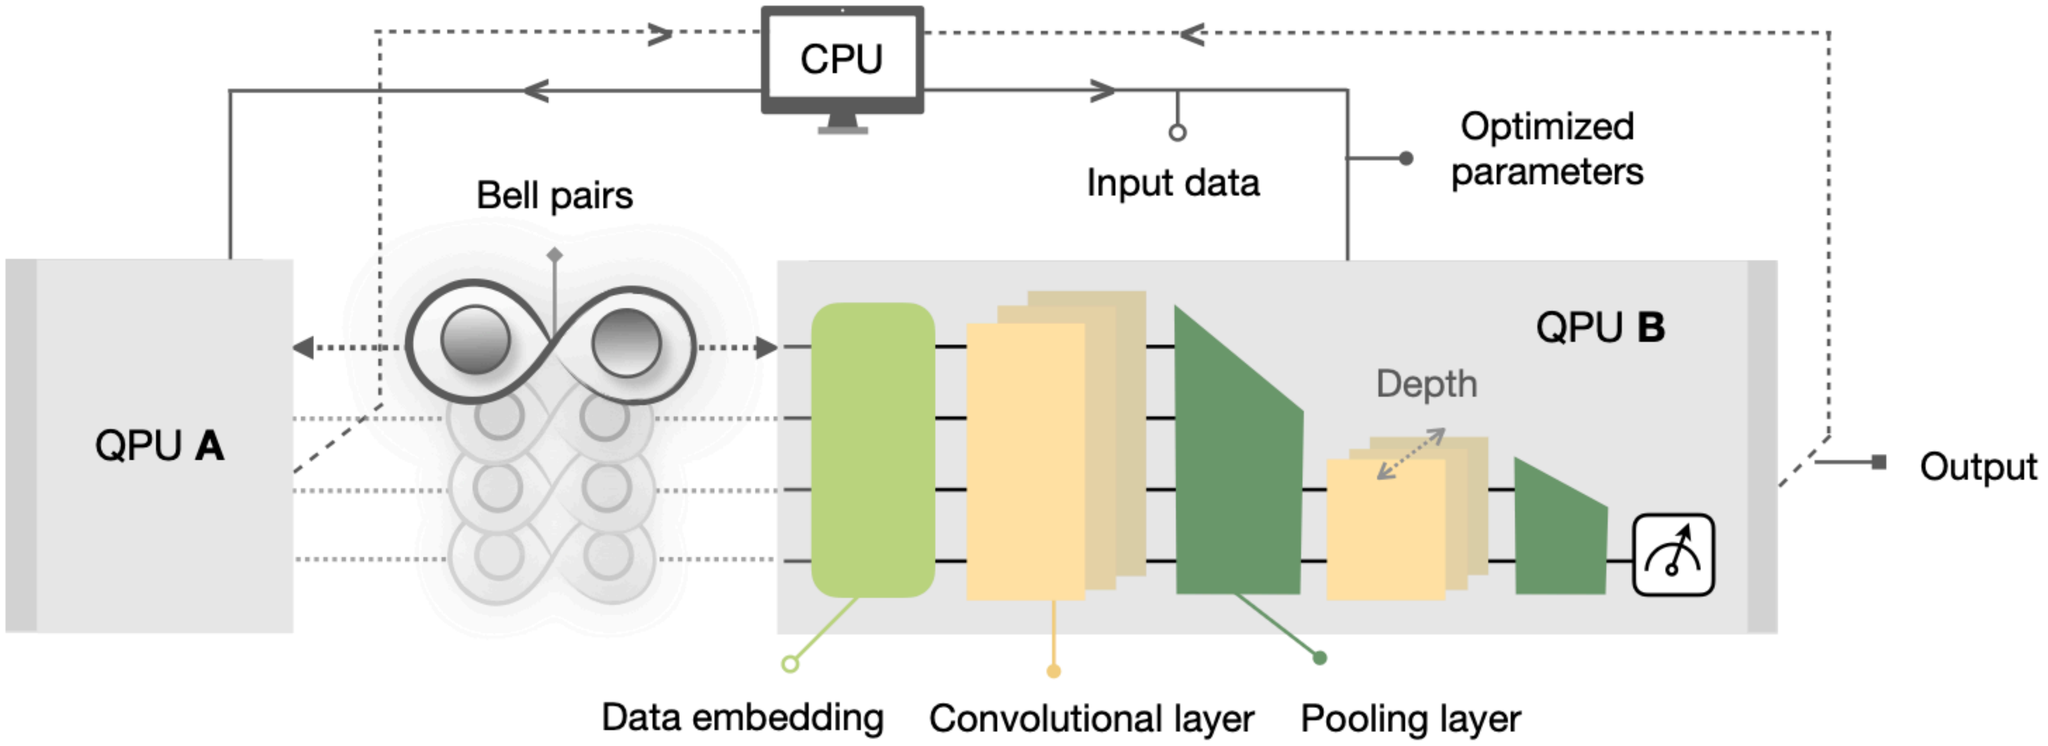

Initial state:

$$|0000\rangle_A|0000\rangle_B \Rightarrow \frac{1}{\sqrt{2}}(|0000\rangle_A|0000\rangle_B + |1000\rangle_A|1000\rangle_B)$$

Results

For two of 4-qubit processors,

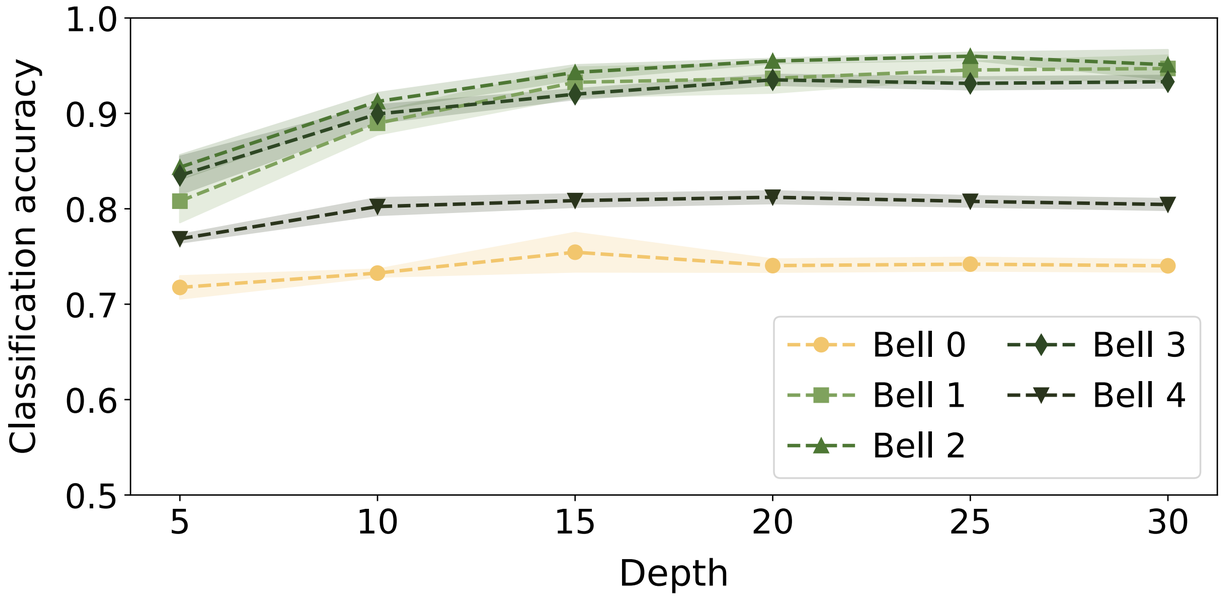

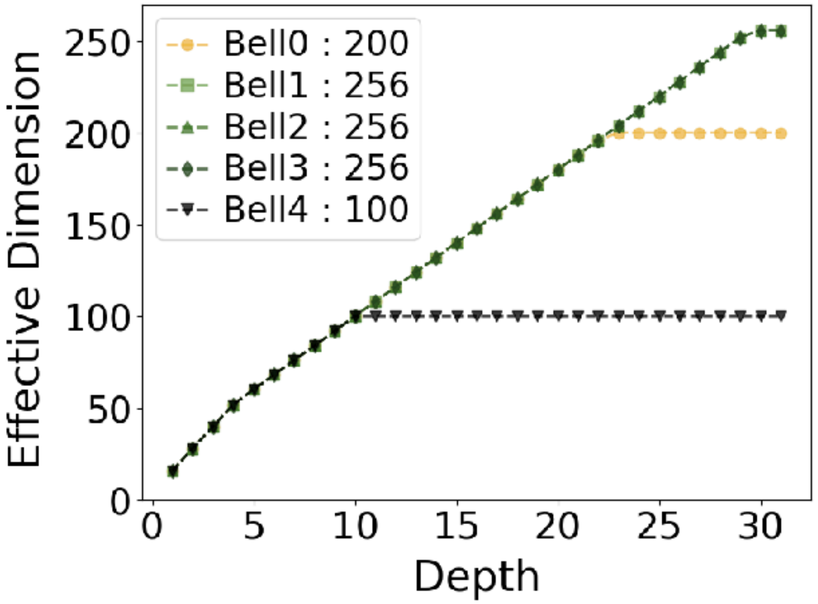

An excess of Bell sources reduces the available parameter (Hilbert) space, thereby lowering learning performance.

Concurrence vs Classification accuracy

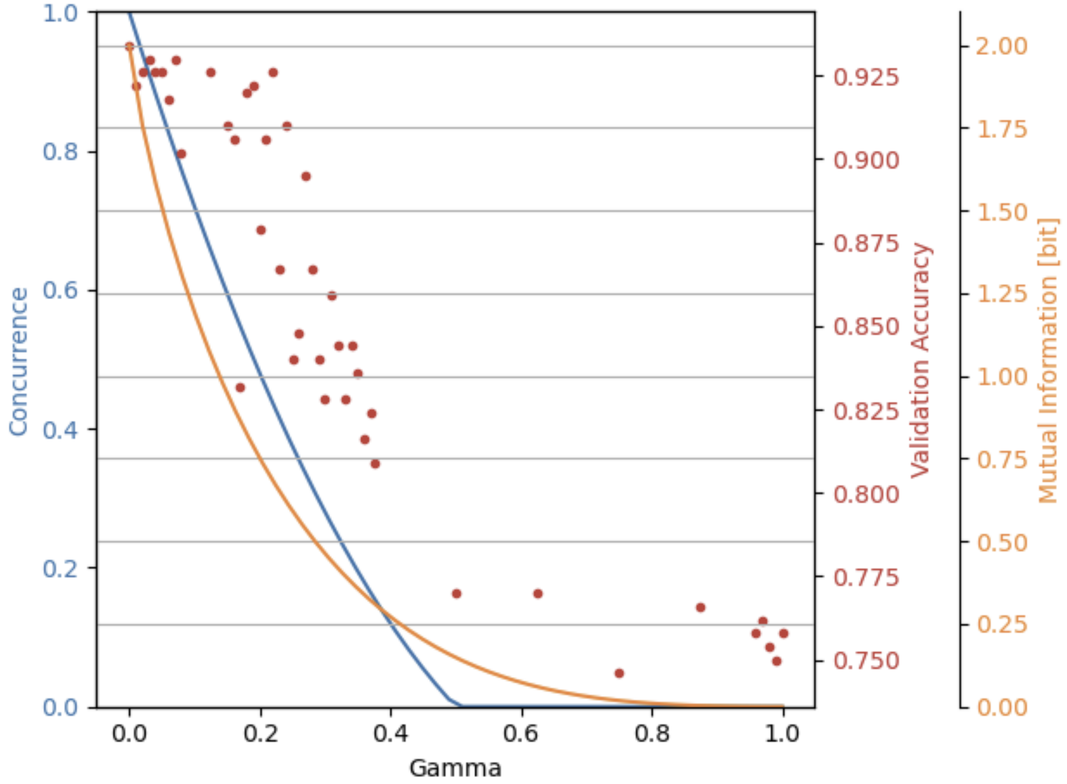

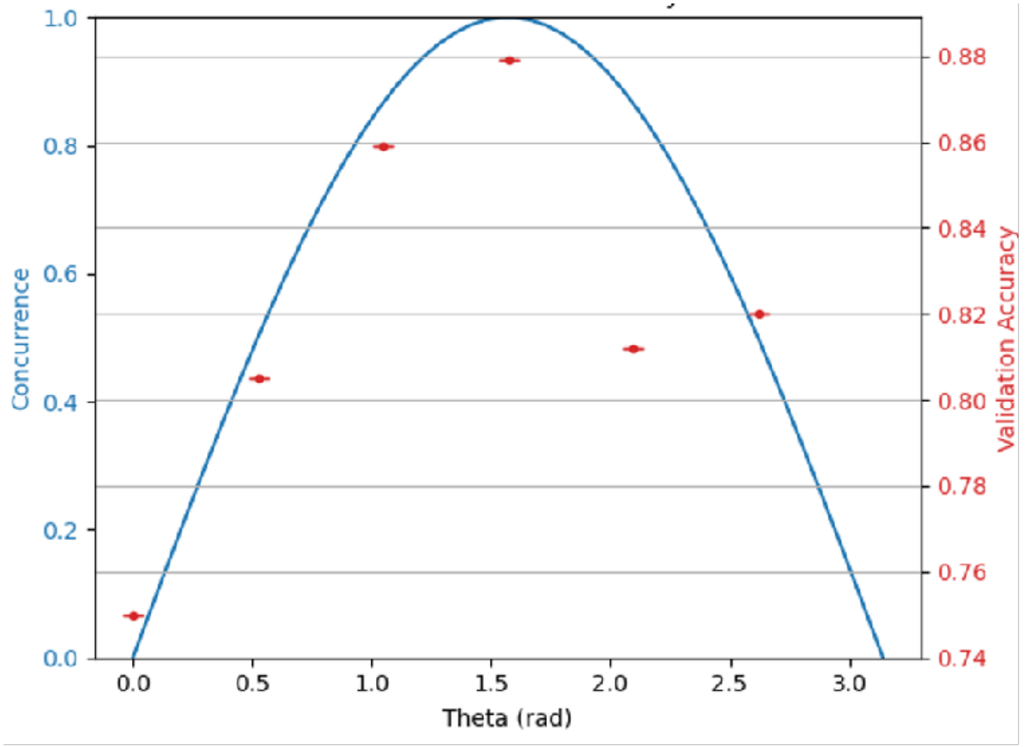

Classification accuracy is strongly correlated with concurrence

Summary

Communication resources

1. *Qubits*
2. *Classical bits*
3. *Ebits (Bell sources)*

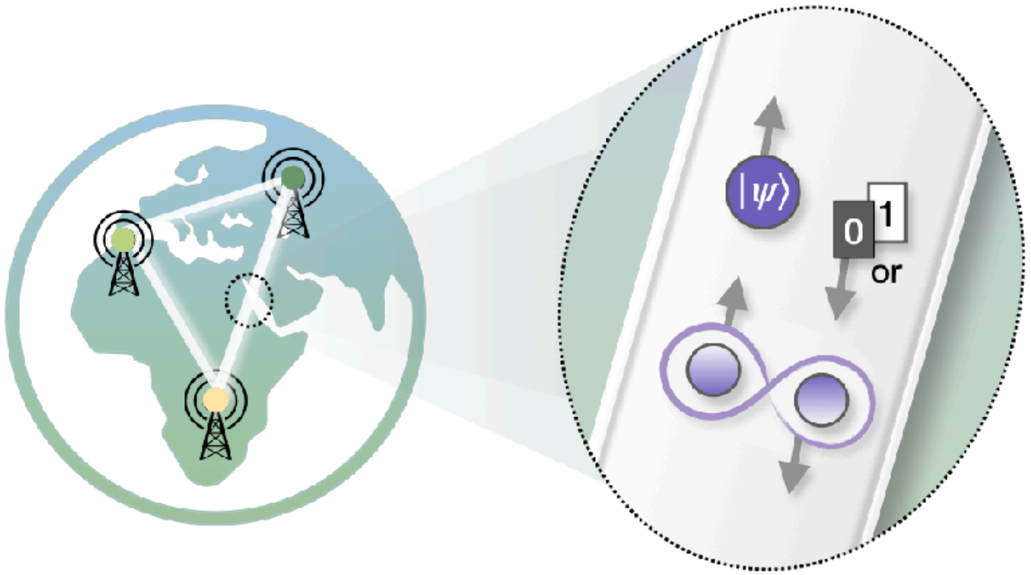

1. Quantum communication (teleportation) is ideal for distributed quantum computing,
   but classical communication and pre-shared entanglement can be comparably effective.

2. In the NISQ era, such practical resources provide a pathway to large-scale, useful applications.

3. Deeper insight into communication resources is essential
   for developing practical distributed quantum computing applications.