Title: Does provable absence of barren plateaus imply classical simulability?

Speakers: Zoë Holmes

Collection/Series: Special Seminars

Subject: Other

Date: March 28, 2025 - 9:00 AM

URL: https://pirsa.org/25030071

Abstract:

A large amount of effort has recently been put into understanding the barren plateau phenomenon. In this perspective talk, we face the increasingly loud elephant in the room and ask a question that has been hinted at by many but not explicitly addressed: Can the structure that allows one to avoid barren plateaus also be leveraged to efficiently simulate the loss classically? We present a case-by-case argument that commonly used models with provable absence of barren plateaus are also in a sense classically simulable, provided that one can collect some classical data from quantum devices during an initial data acquisition phase. This follows from the observation that barren plateaus result from a curse of dimensionality, and that current approaches for solving them end up encoding the problem into some small, classically simulable, subspaces. We end by discussing caveats in our arguments including the limitations of average case arguments, the role of smart initializations, models that fall outside our assumptions, the potential for provably superpolynomial advantages and the possibility that, once larger devices become available, parametrized quantum circuits could heuristically outperform our analytic expectations.

Does provable absence of barren plateaus imply classical simulability?

A genuine question

Zoë Holmes

Many fun arguments were had in the making of this work

arXiv preprint arXiv:2312.09121

Marco Cerezo
LANL

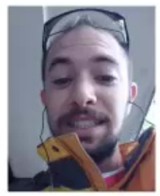

Martin Larocca
LANL

Diego Garcia-Martin
LANL

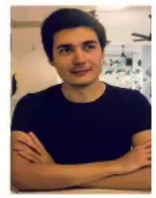

Nahuel Diaz
LANL / La Plata University

Paolo Braccia
LANL

Enrico Fontana
Strathclyde University

Manuel Rudolph
EPFL

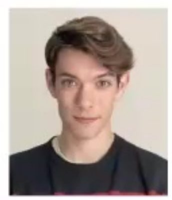

Pablo Bermejo
LANL / Donostia Physics Center

Aroosa Ijaz
LANL / Vector Institute

Supanut Thanasilp
EPFL

Eric Anschuetz
Caltech

Hybrid Variational Quantum Computing

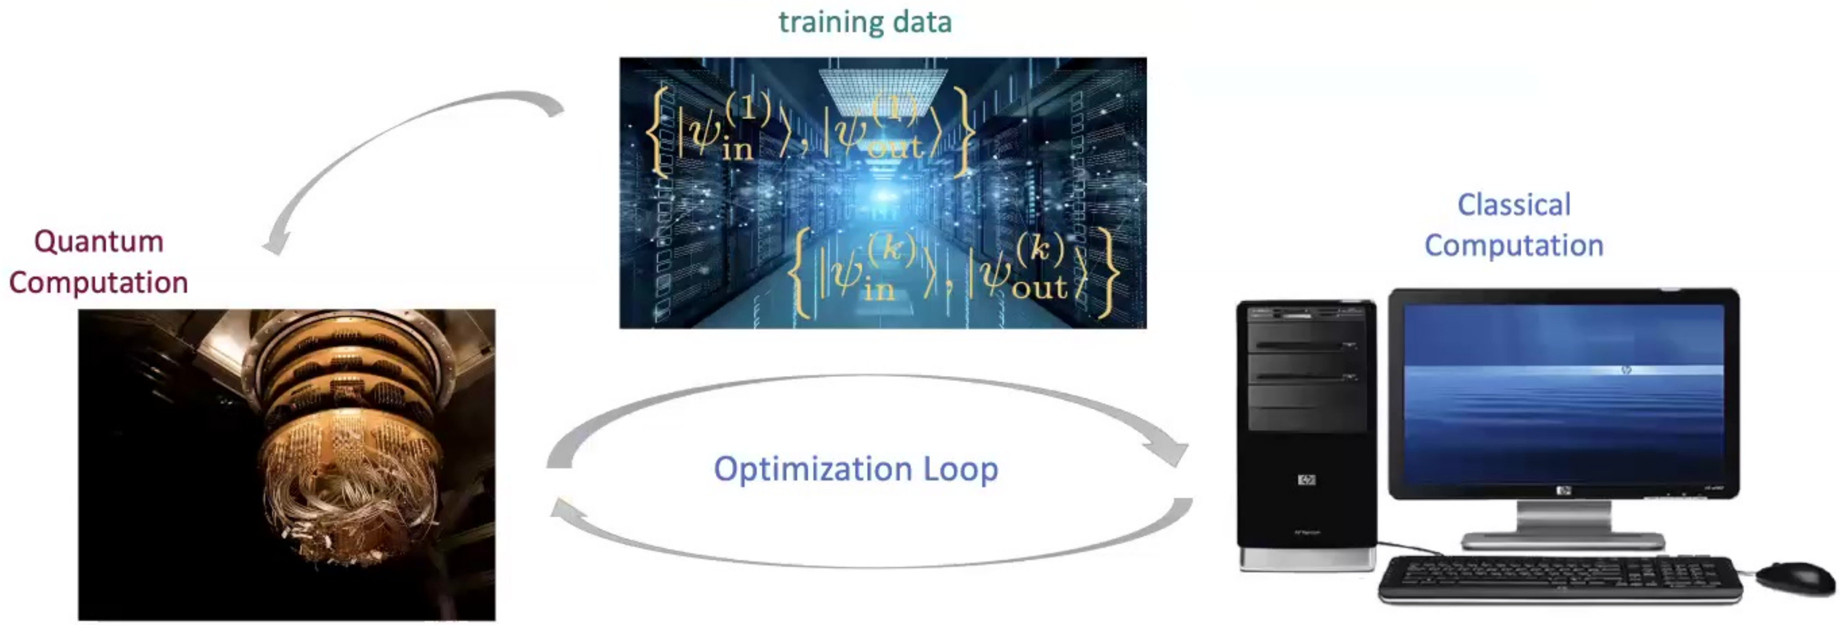

For the purpose of this talk I will focus on:

"Any method that optimizes a parameterized quantum circuit (PQC) by minimizing a quantum cost function (potentially using training data)"

Encapsulates: variational quantum algorithms (VQAs) + many QML methods

Goal: Train a parameterized quantum circuit to minimize a problem-specific loss

Pick faithful loss + PQC s.t. if successfully trained: the optimal parameters/circuit/cost = approx. solution to problem

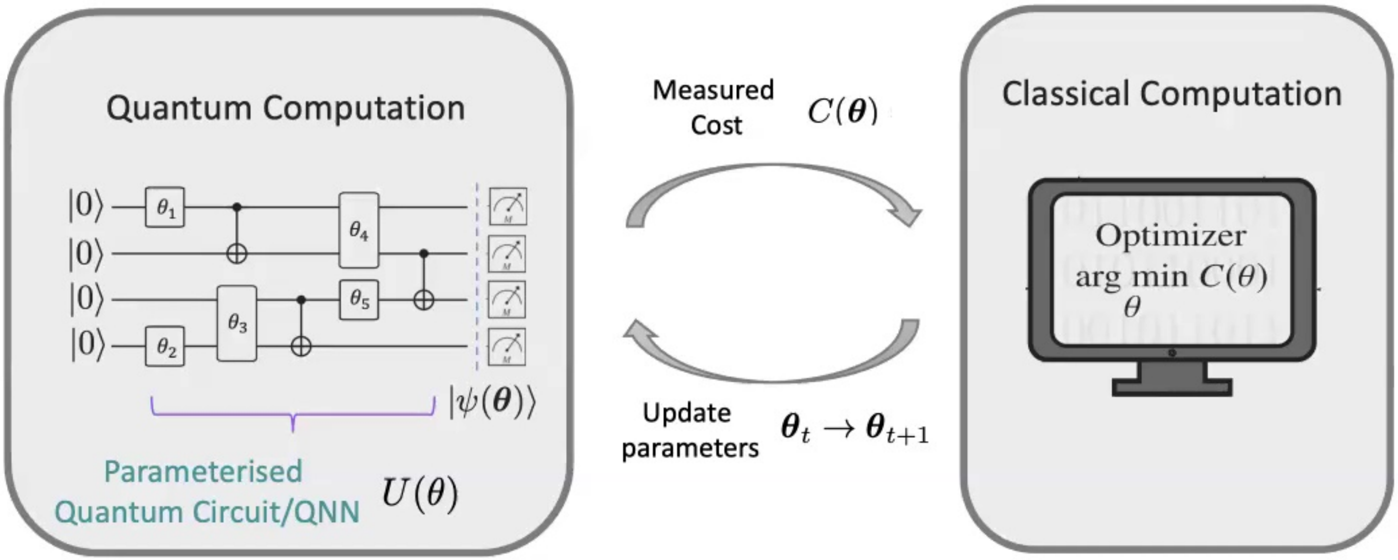

How do you train? Using a hybrid-quantum classical optimization loop

Popular flexible framework

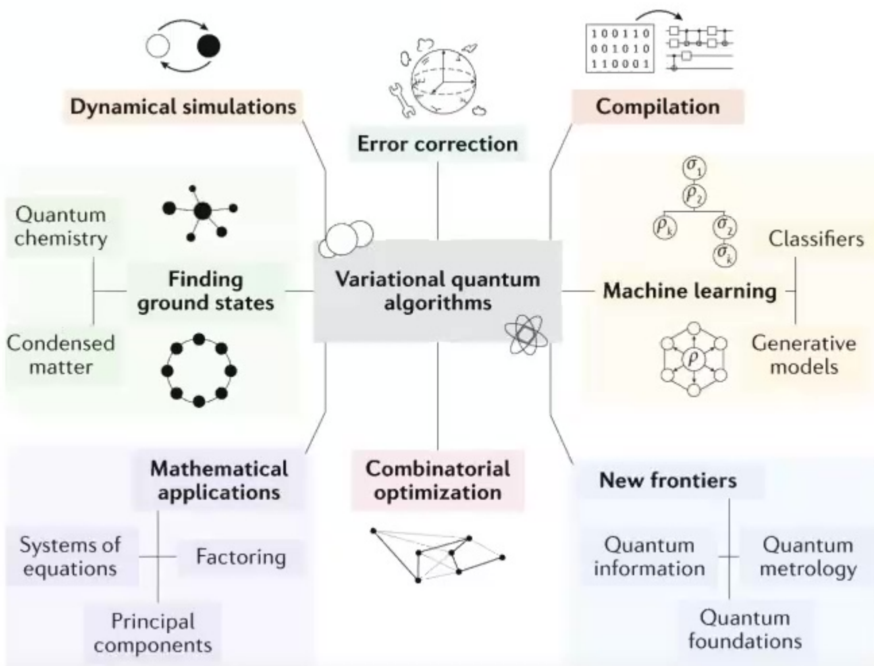

Google Scholar "variational quantum algorithm"

Articles About 1'450 results (**0.10** sec)

Google Scholar "quantum machine learning"

Articles About 17'000 results (**0.05** sec)

Review: *Variational Quantum Algorithms*, **MC** et al, Nature Reviews Physics 3, 625-644 (2021)
(~2000 citations in 3 years)

Ingredients to trainability

1. Expressibility

~ how likely it is that the parameterised unitary contains a solution

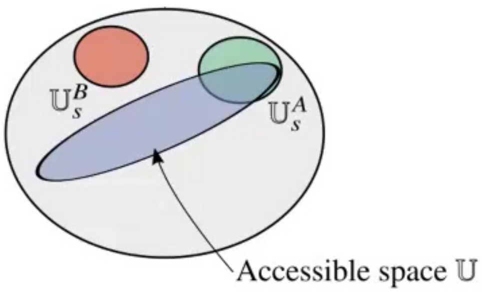

2. Loss gradients (/differences)

~ how easy it is to find a cost minimizing direction

3. Local minima

~ potential to get trapped in poor minima

Cost gradients and shot noise

The cost landscape needs to be sufficiently featured to enable training

Small gradients → high precision required to find loss minimizing direction → resource intensive

($\sim 1/\sigma^2$ shots are required estimate a loss to precision σ)

Barren plateaus

Barren plateau (BP) phenomena:

$$\mathrm{Var}[\partial_k C] \sim \frac{1}{2^n}$$

$$+$$

$$P(|\partial_k C| \geq \delta) \leq \frac{\mathrm{Var}[\partial_k C]}{\delta^2}$$

$$\downarrow$$

Probability of non-zero gradients vanishes exponentially with problem size.

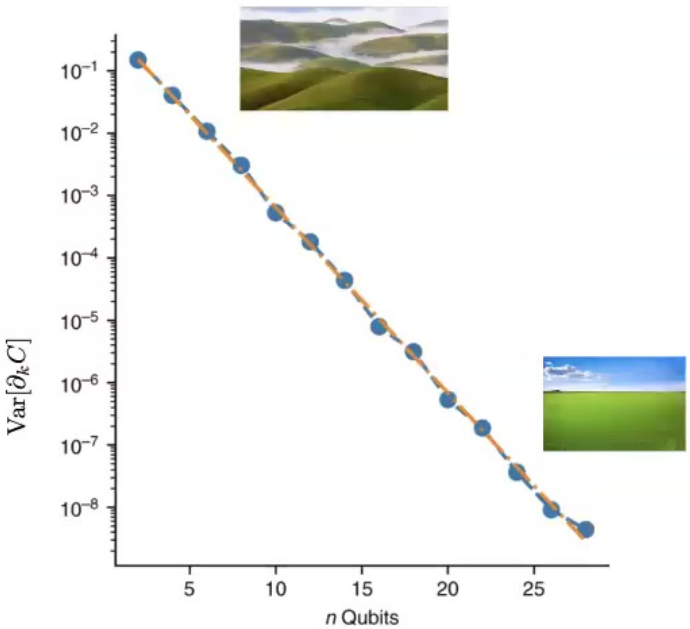

Exponential Cost Concentration

Barren plateau (BP) phenomena:

$$\mathrm{Var}[\cancel{\partial_k} C] \sim \frac{1}{2^n}$$

$$+$$

$$P(|\overset{C-\mu}{\cancel{\partial_k C}}| \geq \delta) \leq \frac{\mathrm{Var}[\cancel{\partial_k} C]}{\delta^2}$$

$$\downarrow$$

Probability of cost deviates from a fixed point vanishes exponentially with problem size.

$$\downarrow$$

Shot required for training grows exponentially with problem size.

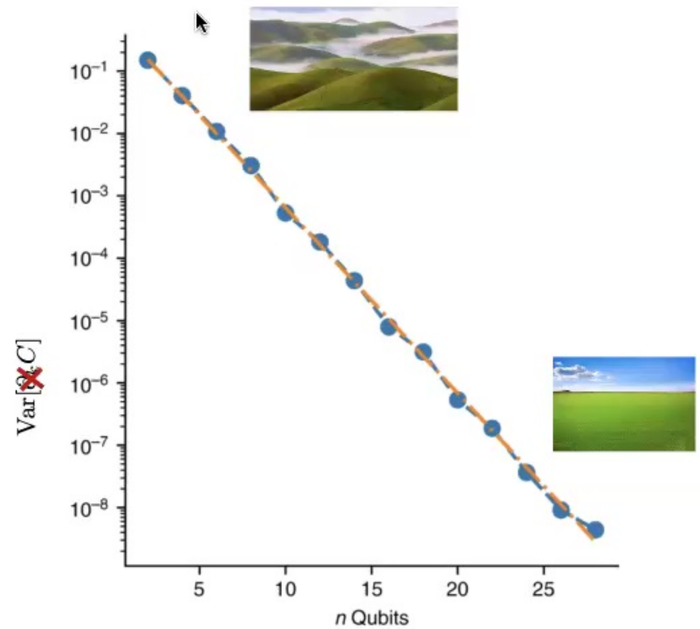

Lots of research on sources of barren plateaus

$$\mathrm{Var}[\partial_k C] \sim \frac{1}{2^n}$$

Choice in circuit

Barren plateaus in quantum neural network training landscapes

Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush & Hartmut Neven

Nature Communications **9**, Article number: 4812 (2018) | Cite this article

- Too expressive

Connecting Ansatz Expressibility to Gradient Magnitudes and Barren Plateaus

Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles
PRX Quantum **3**, 010313 – Published 24 January 2022

Inexpressive

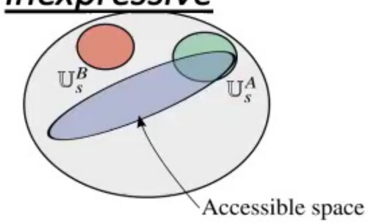

Expressive

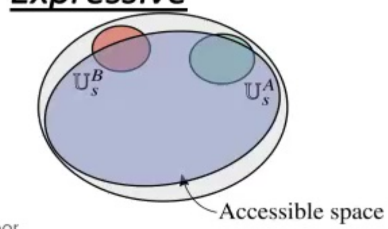

- Too entangling

Entanglement-Induced Barren Plateaus

Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe
PRX Quantum **2**, 040316 – Published 25 October 2021

Choice in target learning problem

Barren Plateaus Preclude Learning Scramblers

Zoë Holmes, Andrew Arrasmith, Bin Yan, Patrick J. Coles, Andreas Albrecht, and Andrew T. Sornbor
Phys. Rev. Lett. **126**, 190501 – Published 12 May 2021

Choice in cost function

Cost function dependent barren plateaus in shallow parametrized quantum circuits

M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio & Patrick J. Coles

Global

$$H = \sigma_1^z \otimes \sigma_2^z \otimes \cdots \otimes \sigma_n^z$$

Local

$$H = \sigma_1^z \otimes \mathbb{I}_2 \otimes \cdots \otimes \mathbb{I}_n$$

Noise

Noise-induced barren plateaus in variational quantum algorithms

Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio & Patrick J. Coles

& lots on provably barren plateau free strategies

$$\mathrm{Var}[\partial_k C] \sim \frac{1}{\mathrm{poly}(n)}$$

No one wants to be (just) the bearer of bad news… certain architectures / methods can be proved to not have BP!

- Shallow hardware-efficient ansatzes with local costs

Cost function dependent barren plateaus in shallow parametrized quantum circuits

M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio & Patrick J. Coles

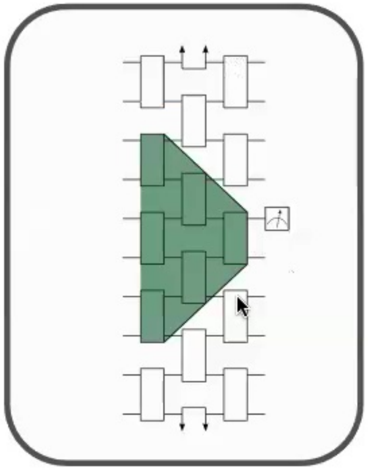

& lots on provably barren plateau free strategies

$$\mathrm{Var}[\partial_k C] \sim \frac{1}{\mathrm{poly}(n)}$$

No one wants to be (just) the bearer of bad news… certain architectures / methods can be proved to not have BP!

- Shallow hardware-efficient ansatzes with local costs

Cost function dependent barren plateaus in shallow parametrized quantum circuits

M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio & Patrick J. Coles

- Quantum Convolutional Neural networks

Absence of Barren Plateaus in Quantum Convolutional Neural Networks

Arthur Pesah, M. Cerezo, Samson Wang, Tyler Volkoff, Andrew T. Sornborger, and Patrick J. Coles
Phys. Rev. X **11**, 041011 – Published 15 October 2021

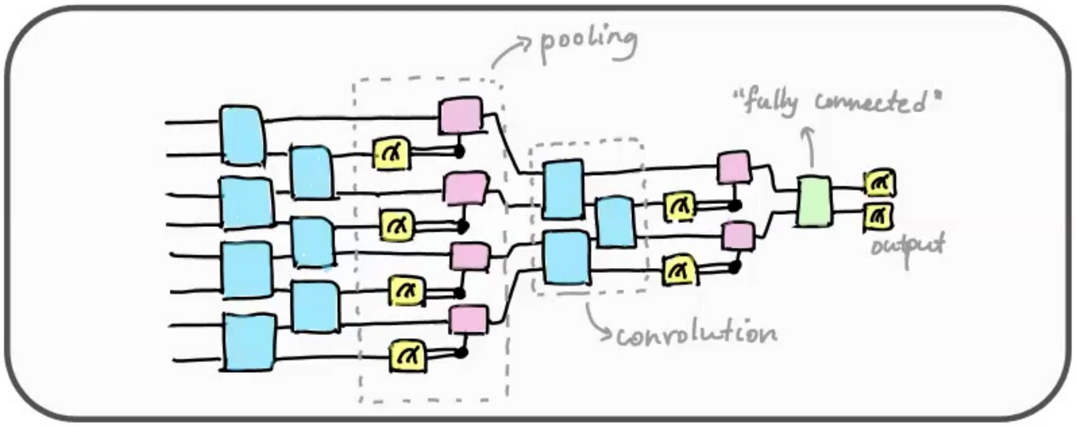

& lots on provably barren plateau free strategies

$$\mathrm{Var}[\partial_k C] \sim \frac{1}{\mathrm{poly}(n)}$$

No one wants to be (just) the bearer of bad news... certain architectures / methods can be proved to not have BP!

- Shallow hardware-efficient ansatzes with local costs
- Quantum Convolutional Neural networks
- Certain Symmetrized Ansatzes
- Small angle initializations

Cost function dependent barren plateaus in shallow parametrized quantum circuits

M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio & Patrick J. Coles

Absence of Barren Plateaus in Quantum Convolutional Neural Networks

Arthur Pesah, M. Cerezo, Samson Wang, Tyler Volkoff, Andrew T. Sornborger, and Patrick J. Coles
Phys. Rev. X **11**, 041011 – Published 15 October 2021

A Unified Theory of Barren Plateaus for Deep Parametrized Quantum Circuits

Michael Ragone,[1,*] Bojko N. Bakalov,[2,*] Frédéric Sauvage,[3] Alexander F. Kemper,[4] Carlos Ortiz Marrero,[5,6] Martín Larocca,[3,7] and M. Cerezo[8,†]

The Adjoint Is All You Need: Characterizing Barren Plateaus in Quantum Ansätze

Enrico Fontana,[1,2] Dylan Herman,[1,*] Shouvanik Chakrabarti,[1] Niraj Kumar,[1] Romina Yalovetzky,[1] Jamie Heredge,[1,3] Shree Hari Sureshbabu,[1] and Marco Pistoia[1]

Escaping from the Barren Plateau via Gaussian Initializations in Deep Variational Quantum Circuits

Kaining Zhang[*], Liu Liu[*], Min-Hsiu Hsieh[†], Dacheng Tao[*]

Trainability Enhancement of Parameterized Quantum Circuits via Reduced-Domain Parameter Initialization

Yabo Wang[1,2], Bo Qi[1,2], Chris Ferrie[3], and Daoyi Dong[4]

Hamiltonian variational ansatz without barren plateaus

Chae-Yeun Park and Nathan Killoran

Barren plateaus = the curse of dimensionality

The loss is expressed as the inner product—a similarity measure

$$\ell_{\boldsymbol{\theta}}(\rho, O) = \mathrm{Tr}\big[U(\boldsymbol{\theta})\rho U^{\dagger}(\boldsymbol{\theta})O\big] = \langle \rho(\boldsymbol{\theta}), O\rangle = \mathrm{Tr}\big[\rho U^{\dagger}(\boldsymbol{\theta})OU(\boldsymbol{\theta})\big] = \langle \rho, O(\boldsymbol{\theta})\rangle,$$

where $\rho(\boldsymbol{\theta}) = U(\boldsymbol{\theta})\rho U^{\dagger}(\boldsymbol{\theta})$; $O(\boldsymbol{\theta}) = U^{\dagger}(\boldsymbol{\theta})OU(\boldsymbol{\theta})$

and $\langle A, B\rangle = \mathrm{Tr}\big[A^{\dagger}B\big]$ the Hilbert-Schmidt inner product.

Red flag:

the inner product between two exponentially large objects will typically be **exponentially small and concentrated**

Claim 0: BPs = Curse of dimensionality.

When does this curse of dimensionality kick in?

$$\ell_{\boldsymbol\theta}(\rho, O) = \mathrm{Tr}\left[U(\boldsymbol\theta)\rho U^\dagger(\boldsymbol\theta) O\right] = \langle \rho, O(\boldsymbol\theta)\rangle$$

Let $O = \sum_\lambda c_\lambda P_\lambda$

Where can O go when we Heisenberg backpropagate it....

Cost function dependent barren plateaus in shallow parametrized quantum circuits

M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio & Patrick J. Coles

Can we explore an exponentially large space?

These questions can be answered by computing $\left|\langle O(\boldsymbol\theta), P_j\rangle\right|^2$ for all elements of a basis of the Hilbert space

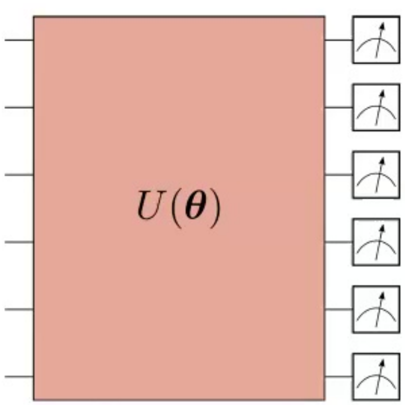

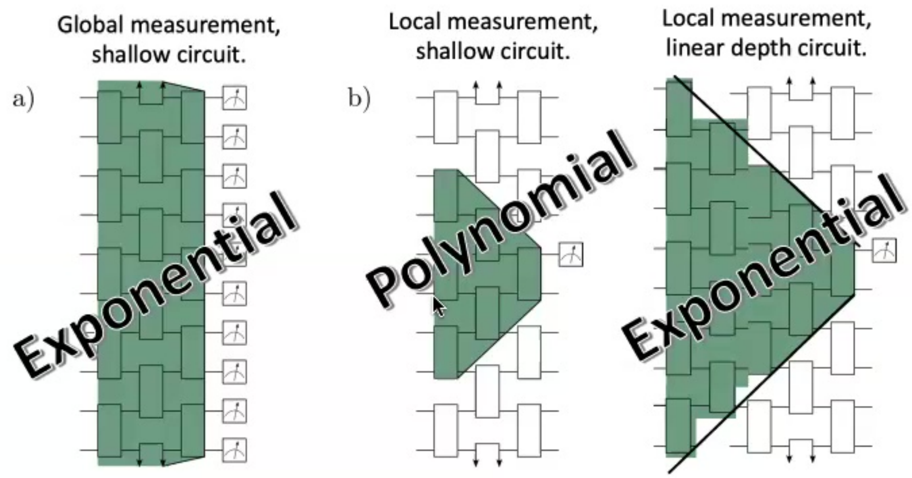

- Circuit depth
- Circuit structure
- Measurement O

determine where we can go.

When does this curse of dimensionality kick in?

$\ell_{\boldsymbol{\theta}}(\rho, O) = \mathrm{Tr}\left[U(\boldsymbol{\theta})\rho U^{\dagger}(\boldsymbol{\theta})O\right] = \langle \rho, O(\boldsymbol{\theta})\rangle$

Take $O = \sum_{\lambda} c_{\lambda} P_{\lambda}$, and define as $\mathcal{B}_{\lambda}$ the subspace associated to each P_{λ} under the adjoint action of $U(\boldsymbol{\theta})$.
We have

$$\ell_{\boldsymbol{\theta}}(\rho, O) = \sum_{\lambda} c_{\lambda} \langle \rho, P_{\lambda}(\boldsymbol{\theta})\rangle = \sum_{\lambda} c_{\lambda} \langle \rho_{\lambda}, P_{\lambda}(\boldsymbol{\theta})\rangle$$

where ρ_{λ} is the projection of ρ onto $\mathcal{B}_{\lambda}$.

The loss is the sum of the inner products in each subspace!

If any of the subspaces is only polynomially large the loss can be non-concentrated

If ρ is too entangled (volume law) the reduced state will be maximally mixed and the loss will be exponentially suppressed.

So, ρ must not be too entangled (satisfy an area law).

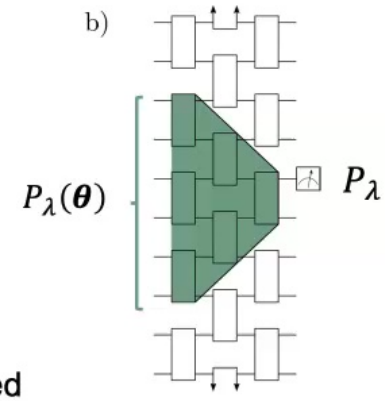

Entanglement-Induced Barren Plateaus

Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe
PRX Quantum **2**, 040316 – Published 25 October 2021

Provable absence of BPs = Problem lives in polynomially subspace

<u>Claim 1:</u> Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

- Shallow hardware-efficient ansatzes with local costs
- Quantum Convolutional Neural networks
- Certain Symmetrized Ansatzes
- Small angle initializations

Operators in Polynomial subspace

log(n) local neighbouring Pauli operators

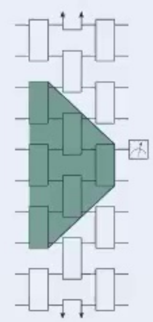

Provable absence of BPs = Problem lives in polynomially subspace

Claim 1: Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

- Shallow hardware-efficient ansatzes with local costs
- Quantum Convolutional Neural networks
- Certain Symmetrized Ansatzes
- Small angle initializations

Operators in the small lie algebra of the generators

Provable absence of BPs = Problem lives in polynomially subspace

In the paper we provide a more detailed analysis…

<u>Claim 1:</u> Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

Problem instance $\mathcal{C}$ based on	Refs.	Operators in the polynomial-sized $\mathcal{B}_\lambda$	O and ρ leading to $\mathcal{C} \subset \overline{\text{BP}}$
Shallow hardware efficient ansatz	[24–27, 29–31]	$\mathcal{O}(\log(n))$ neighboring qubits (P)	Local O, area law ρ
Generic shallow locally circuits	[105, 106][a]	$\mathcal{O}(1)$-weight qubits (E)	Local O, area law ρ
Quantum convolutional neural network	[15]	$\mathcal{O}(1)$-weight qubits (E)	Local O, area law ρ
$U(1)$-equivariant	[22, 50, 51]	Proj. with $\mathcal{O}(1)$ Hamming weight (P)	Equivariant O, $\rho \in \mathcal{B}_\lambda$
S_n-equivariant	[67]	Permutation equivariant (P)	$O \in \mathcal{B}_\lambda$, $\text{Tr}[\rho_\lambda^2] \in \Omega(1/\text{poly}(n))$
Matchgate circuit	[54]	Product of $\mathcal{O}(1)$ Majoranas (P)	$O \in \mathcal{B}_\lambda$, $\text{Tr}[\rho_\lambda^2] \in \Omega(1/\text{poly}(n))$
Small angle initialization	[39, 55, 61]	$\mathcal{O}(\text{poly}(n))$ Pauli Operators (E)	Local O, area law ρ
Small Lie algebra $\mathfrak{g}$	[52, 58, 65]	Operators in $\mathfrak{g}$ (P)	$O \in i\mathfrak{g}$, $\text{Tr}[\rho_{\mathfrak{g}}^2] \in \Omega(1/\text{poly}(n))$
Non-unital noisy-circuits	[70]	$\mathcal{O}(\log(n))$ qubits (E)	Local O, any ρ
Dynamic circuits	[73]	$\mathcal{O}(1)$-weight qubits (E)	Local O, any ρ
Quantum generative modeling[b]	[29, 32]	Tensor networks (e.g., MPS) (P)	O, ρ computational basis proj.

Crucially: very proof of absence of BPs allowed us to find poly subspaces (more on this later)

Classical simulability Definitions

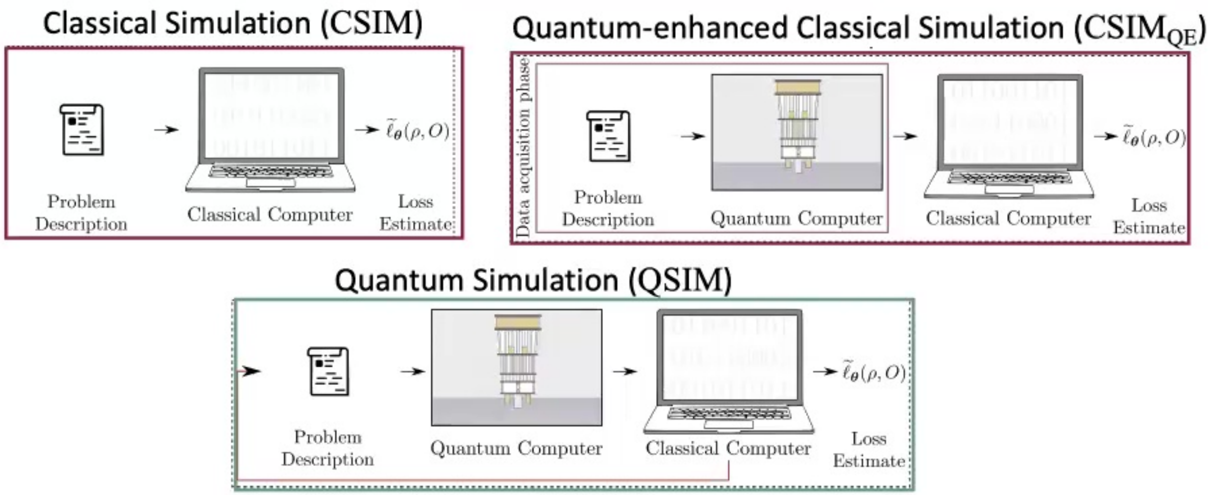

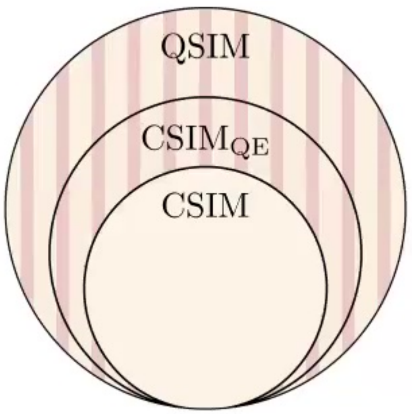

A quantum advantage is possible for any problem in $CSIM_{QE} \cap \neg CSIM$ or $QSIM \cap \neg CSIM$

An advantage from adaptively running a parameterized quantum circuit is only possible for problems in $QSIM \cap \neg CSIM_{QE}$

Provable absence of BPs = Problem lives in polynomially subspace

Back to our case-by-case analysis:

Claim 1: Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

In each of these cases.... $U(\boldsymbol{\theta})$ can be classically simulated
& so the loss can be classically simulated for classical initial states/measurements

Shallow hardware efficient ansatz
Quantum convolutional neural network
$U(1)$-equivariant
S_n-equivariant
Matchgate circuit
Small angle initialization
Small Lie algebra $\mathfrak{g}$
Quantum generative modeling[a]

(Nothing really new here - this was increasingly discussed but hadn't fully sunk in all corners)

Provable absence of BPs = Problem lives in polynomially subspace

Case-by-case analysis:

<u>Claim 1:</u> Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

In each of these cases…. $U(\boldsymbol{\theta})$ can be classically simulated
& so the loss can be classically simulated for classical initial states/measurements

Tensor networks

Shallow hardware efficient ansatz
Quantum convolutional neural network
$U(1)$-equivariant
S_n-equivariant
Matchgate circuit
Small angle initialization
Small Lie algebra $\mathfrak{g}$
Quantum generative modeling[a]

Provable absence of BPs = Problem lives in polynomially subspace

Case-by-case analysis:

<u>Claim 1:</u> Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

In each of these cases…. $U(\boldsymbol{\theta})$ can be classically simulated
& so the loss can be classically simulated for classical initial states/measurements

	Shallow hardware efficient ansatz
	Quantum convolutional neural network
(Particle number	conserving) $U(1)$-equivariant
(Permutation	invariant) S_n-equivariant
	Matchgate circuit
	Small angle initialization
	Small Lie algebra $\mathfrak{g}$
	Quantum generative modeling[a]

Gsim

Efficient Solvability of Hamiltonians and Limits on the Power of Some Quantum Computational Models

Rolando Somma,[1] Howard Barnum,[1] Gerardo Ortiz,[1] and Emanuel Knill[2]

Efficient classical algorithms for simulating symmetric quantum systems

Eric R. Anschuetz[1], Andreas Bauer[2], Bobak T. Kiani[3], and Seth Lloyd[4,5]

Fermion Sampling: A Robust Quantum Computational Advantage Scheme Using Fermionic Linear Optics and Magic Input States

Michał Oszmaniec, Ninnat Dangniam, Mauro E.S. Morales, and Zoltán Zimborás
PRX Quantum 3, 020328 – Published 9 May 2022

Classical and Quantum Algorithms for Orthogonal Neural Networks

Iordanis Kerenidis[1,2], Jonas Landman[1,2], and Natansh Mathur[3,2]

Lie-algebraic classical simulations for variational quantum computing

Matthew L. Goh,[1,2,*] Martin Larocca,[1,3] Lukasz Cincio,[1] M. Cerezo,[4] and Frédéric Sauvage[1]

Provable absence of BPs = Problem lives in polynomially subspace

Case-by-case analysis:

Claim 1: Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

In each of these cases…. $U(\boldsymbol{\theta})$ can be classically simulated
& so the loss can be classically simulated for classical initial states/measurements

Shallow hardware efficient ansatz
Quantum convolutional neural network
$U(1)$-equivariant
S_n-equivariant
Matchgate circuit
Small angle initialization
Small Lie algebra $\mathfrak{g}$
Quantum generative modeling[a]

But I used to say: "for non-classically simulable initial states/measurements
we would need to run the VQA on the quantum computer"

Classically simulating provably BP free losses

Recall from a couple of slides back that we can:

Take $O = \sum_\lambda c_\lambda P_\lambda$, and define as $\mathcal{B}_\lambda$ the subspace associated to each P_λ under the adjoint action of $U(\boldsymbol{\theta})$. We have

$$\ell_{\boldsymbol{\theta}}(\rho, O) = \sum_\lambda c_\lambda \langle \rho_\lambda, P_\lambda(\boldsymbol{\theta}) \rangle$$

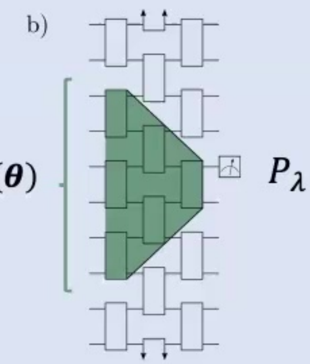

where ρ_λ is the projection of ρ onto $\mathcal{B}_\lambda$.
The loss is the sum of the inner products in each subspace!

So to simulate this loss we just need to compute a basis of $\mathcal{B}_\lambda$ and project ρ_λ onto that basis.

If we do not have a classical representation of ρ_λ then we can do this on a quantum computer.
& crucially, for provably BP free architectures live in polynomial subspaces… so this only takes a polynomial number of measurements…. i.e. the problem is in $\boldsymbol{CSIM_{QE}}$

Classically simulating provably BP free losses

If we do a case-by-case analysis of BP-free architectures/methods used in the literature, we can see that:

Claim 2: Problems in known polynomial subspaces are classically simulable (potentially requiring data from a quantum computer).

Problem instance $\mathcal{C}$ based on	Tomographic procedure for ρ	Simulation algorithm based on
Shallow hardware efficient ansatz	Pauli classical shadows [94]	Light-cone sim. reduced $U(\boldsymbol{\theta})$
Generic shallow locally circuits	Pauli classical shadows	Pauli Propagation [79]
Quantum convolutional neural network[a]	Pauli classical shadows	Pauli Propagation
$U(1)$-equivariant	Computational basis measurement	Givens Rotations [110]
S_n-equivariant	Permutation invariant shadows [111]	$\mathfrak{g}$-sim [112]
Matchgate circuit	Expectation value of Pauli operators	$\mathfrak{g}$-sim
Small angle initialization	Pauli measurements	Tensor Networks [82], Pauli Prop. [80]
Small Lie algebra $\mathfrak{g}$	Expectation value of algebra elements	$\mathfrak{g}$-sim

Punchline: In none of the standard problem instances with provable absence of barren plateaus does the parametrized quantum circuit need to be implemented on a quantum computer in order to estimate the loss in polynomial time.

Where does this leave us?

This must be wrong.

Do we want to kill the field?

I'm gonna have to do fault-tolerant now…

There are caveats!

So many new questions!

This means a better understanding + new opportunities for the field!

Caveats: Counterexample

Case-by-case, not absolute, argument

I. We can construct examples of non-concentrated losses that are not classically simulable based on cryptographic hardness, e.g., we smuggle in the discrete logarithm problem

II. These examples do not resemble current mainstream variational quantum algorithms.

III. Break assumption that comparing objects in exponentially large spaces leads to concentrated expectation values as the circuits are structured

On the relation between trainability and dequantization of variational quantum learning models

Elies Gil-Fuster,[1,2] Casper Gyurik,[3] Adrián Pérez-Salinas,[3,4] and Vedran Dunjko[3,5]

Caveats: Counterexample

Case-by-case, not absolute, argument

I. We can construct examples of non-concentrated losses that are not classically simulable based on cryptographic hardness, e.g., we smuggle in the discrete logarithm problem

II. These examples do not resemble current mainstream variational quantum algorithms.

III. Break assumption that comparing objects in exponentially large spaces leads to concentrated expectation values as the circuits are structured

In its current form I do not find this caveat so interesting… but

Caveats: Special initializations

Both barren plateaus and (in places) our notion of simulability are average case notions

A barren plateaus can have substantial gradients in exp. small subregion

In some cases, simulation might also not be possible in an exp. small subregion

Potential offered by warm starts?

A unifying account of warm start guarantees for patches of quantum landscapes

Hela Mhiri,[1,2,*] Ricard Puig,[1,*] Sacha Lerch,[1] Manuel S. Rudolph,[1]
Thiparat Chotibut,[3] Supanut Thanasilp,[1,3] and Zoë Holmes[1]
[1]*Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*
[2]*Laboratoire d'Informatique de Paris 6, CNRS, Sorbonne Universite, 4 Place Jussieu, 75005 Paris, France*
[3]*Chula Intelligent and Complex Systems, Department of Physics,*
Faculty of Science, Chulalongkorn University, Bangkok, Thailand, 10330

Caveats: Special initializations

A unifying account of warm start guarantees for patches of quantum landscapes

Hela Mhiri,[1,2,*] Ricard Puig,[1,*] Sacha Lerch,[1] Manuel S. Rudolph,[1]
Thiparat Chotibut,[3] Supanut Thanasilp,[1,3] and Zoë Holmes[1]
[1]Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
[2]Laboratoire d'Informatique de Paris 6, CNRS, Sorbonne Université, 4 Place Jussieu, 75005 Paris, France
[3]Chula Intelligent and Complex Systems, Department of Physics,
Faculty of Science, Chulalongkorn University, Bangkok, Thailand, 10330

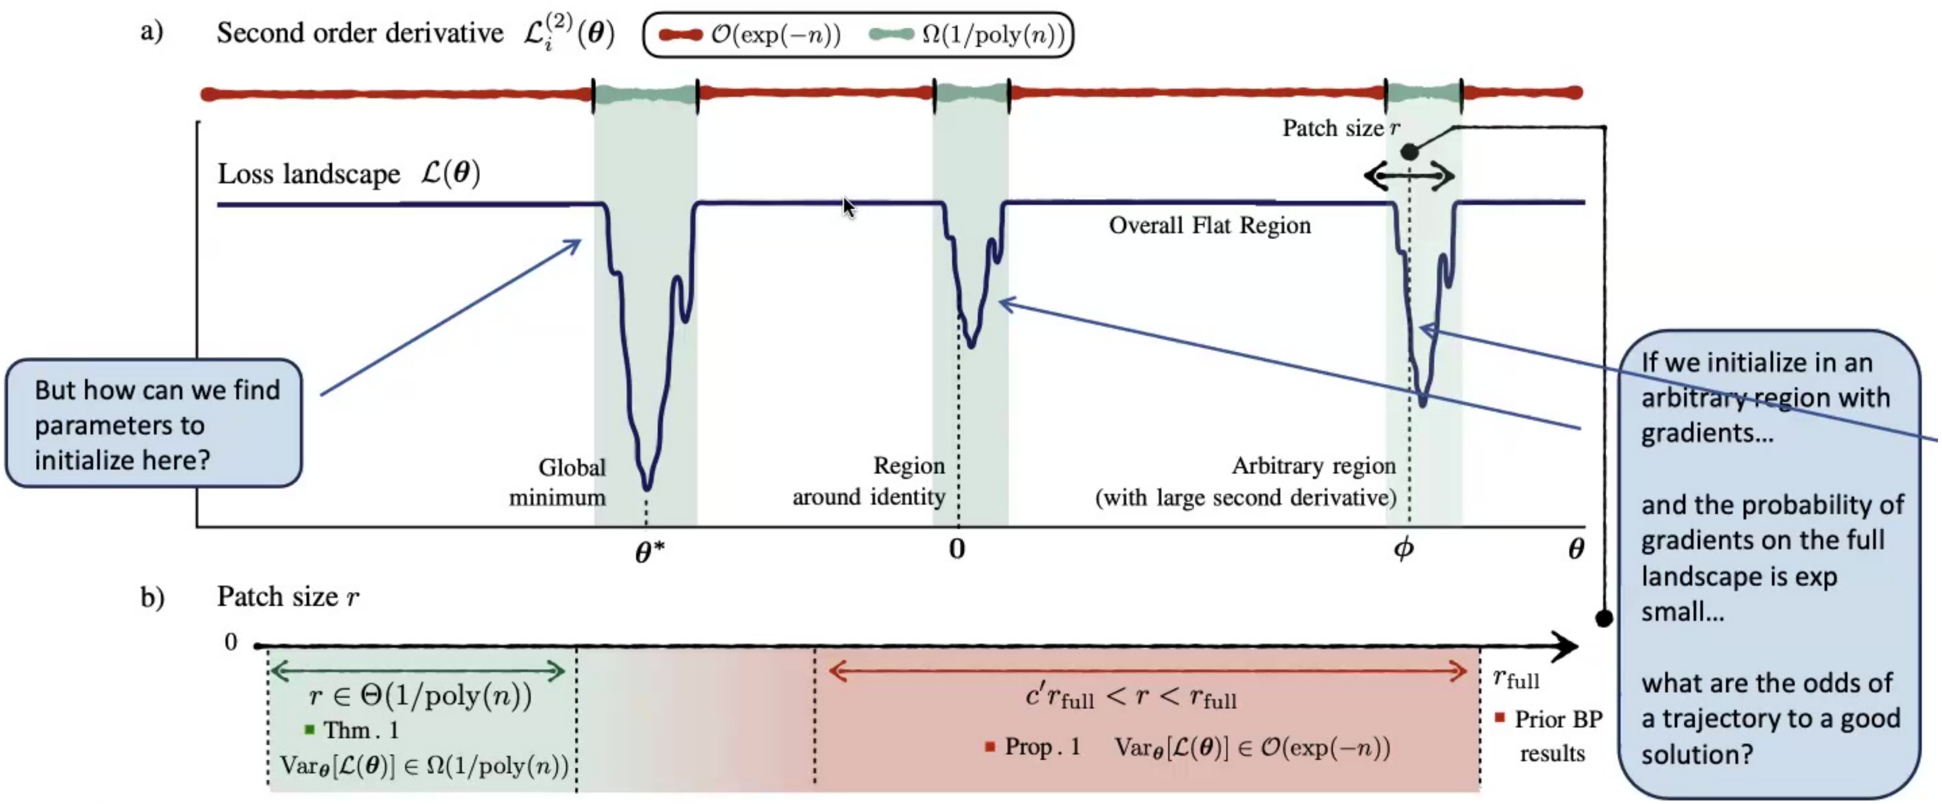

Caveats: Special initializations

Warm starts are 'consistent' with our guarantees…. i.e., these cases are also technically in $\boldsymbol{CSIM_{QE}}$

A unifying account of warm start guarantees for patches of quantum landscapes

Hela Mhiri,[1,2,*] Ricard Puig,[1,*] Sacha Lerch,[1] Manuel S. Rudolph,[1]
Thiparat Chotibut,[3] Supanut Thanasilp,[1,3] and Zoë Holmes[1]

[1] *Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*
[2] *Laboratoire d'Informatique de Paris 6, CNRS, Sorbonne Universite, 4 Place Jussieu, 75005 Paris, France*
[3] *Chula Intelligent and Complex Systems, Department of Physics, Faculty of Science, Chulalongkorn University, Bangkok, Thailand, 10330*

← Can prove absence of barren plateau guarantees

Efficient quantum-enhanced classical simulation for patches of quantum landscapes

Sacha Lerch,[1,*] Ricard Puig,[1,*] Manuel S. Rudolph,[1,*]
Armando Angrisani,[1] Tyson Jones,[1] M. Cerezo,[2,3] Supanut Thanasilp,[1,4] and Zoë Holmes[1]

← Can classically *surrogate* those regions using hardware data

But they push beyond the spirit of 'classically simulable after collecting data from quantum device' because they can require significant quantum resources in the data collection phase….

Caveats: not a strict dequantization!

The quantum computer can still be needed to collect data! (If initial state/measurement is quantum or if using a warm start in a high entanglement / high magic region of the landscape)

But it is used non-adaptively to create a surrogate which is then used for training

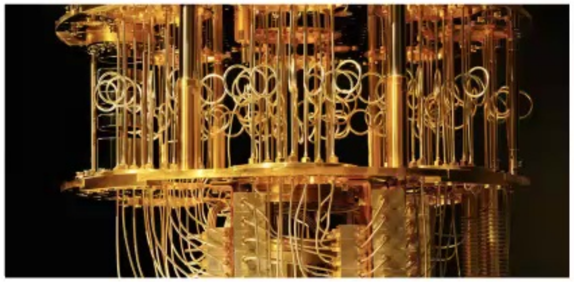

In some cases all you need to do is collect a simple classical shadow… in which case you might not even need a universal quantum computer!

Caveats: not a strict dequantization!

In some cases all you need to do is collect a simple classical shadow… in which case you might not even need a universal quantum computer!

Quantum Convolutional Neural Networks are (Effectively) Classically Simulable

Pablo Bermejo,[1,2,3] Paolo Braccia,[4] Manuel S. Rudolph,[5] Zoë Holmes,[5] Lukasz Cincio,[4] and M. Cerezo[1,*]

$$H = -J\sum_{i=1}^{n-2} Z_i X_{i+1} Z_{i+2} - h_1 \sum_{i=1}^{n} X_i - h_2 \sum_{i=1}^{n-1} X_i X_{i+1}$$

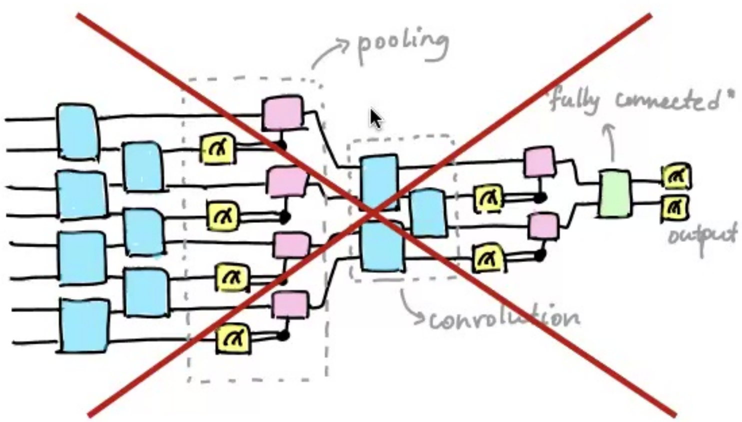

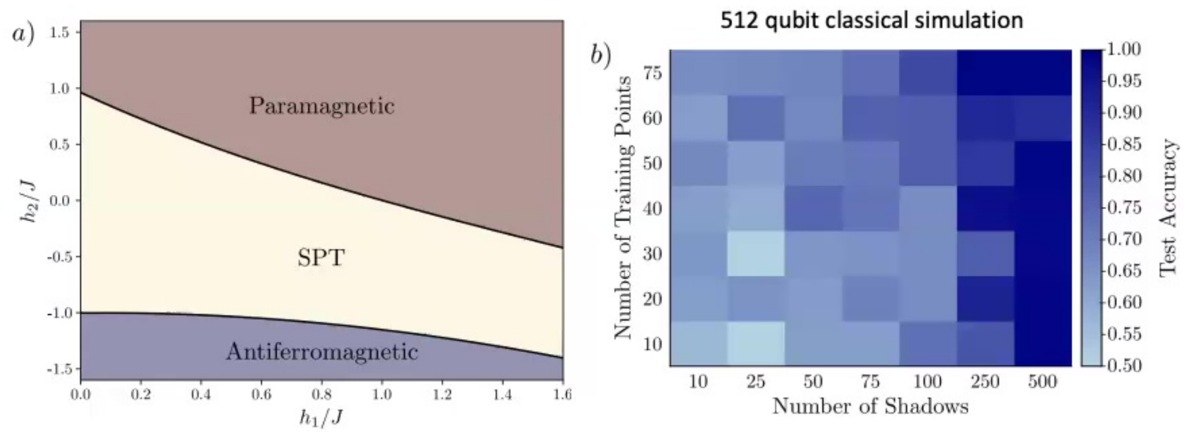

Caveats: not a strict dequantization!

But in other cases (e.g. warm starts away from Clifford / low entanglement circuits)…

i. Complex circuits may still need to be run on hardware

ii. The advantages/disadvantages of surrogating become blurred

Efficient quantum-enhanced classical simulation for patches of quantum landscapes

Sacha Lerch,[1,*] Ricard Puig,[1,*] Manuel S. Rudolph,[1,*] Armando Angrisani,[1] Tyson Jones,[1] M. Cerezo,[2,3] Supanut Thanasilp,[1,4] and Zoë Holmes[1]

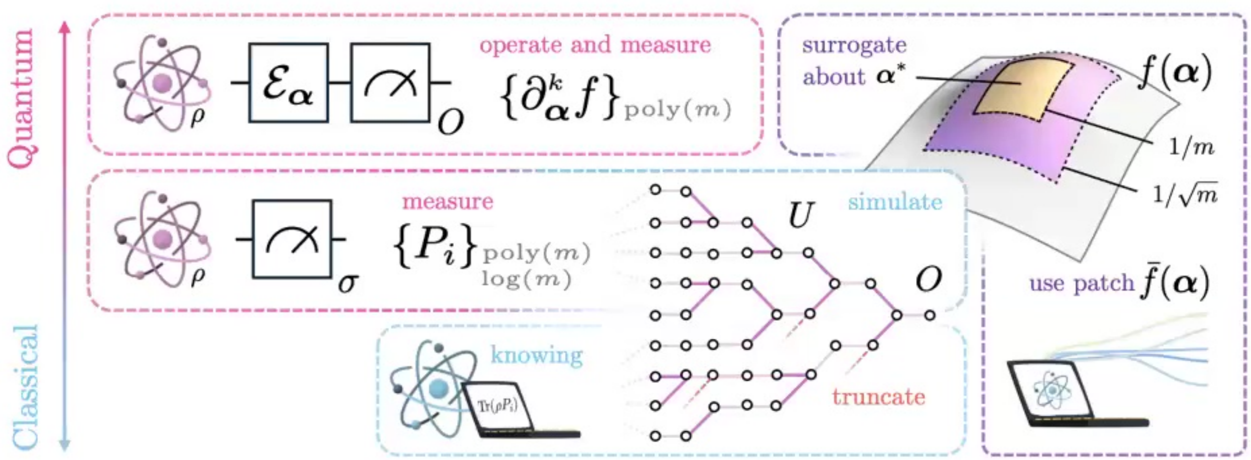

Caveats: Who cares about proofs

We rely on proofs of absence of BPs.

Could heuristically find large gradients but no identifiable poly-subspace?

Analogous to classical case?

Remember precision is much more expensive than classically

& The fact classical ML is so successful means we have a high bar to beat.

Conclusions

Claim 0: BPs = Curse of dimensionality.

Claim 1: Provably barren plateau-free architectures live in classically identifiable polynomial subspaces.

Claim 2: Problems in known polynomial subspaces are classically simulable (potentially requiring data from a quantum computer).

Lots of caveats / future opportunities