Characterising, simulating and boosting quantum circuits with classical neural nets

Leandro Aolita

QRC-TII, Abu Dhabi & IF-UFRJ, Rio de Janeiro

Impressive progress on experimental quantum computations and simulations

Micro-fabricated trapped-ion architectures

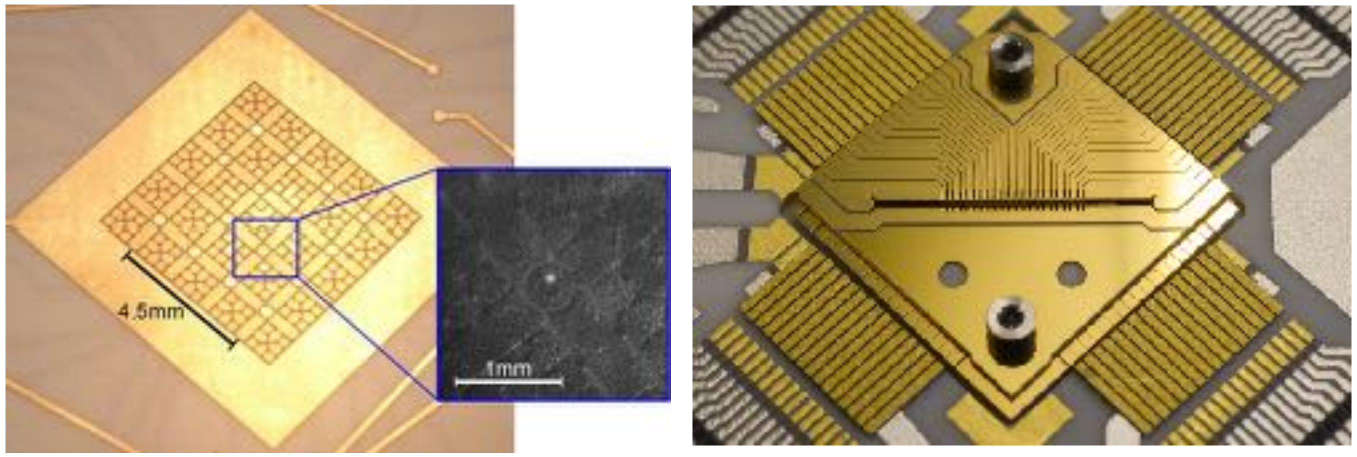

Wineland; Blatt; Monroe; Leibfried; Roos; Häffner; Schmidtkaler; Schätz; Kim; etc.

On-chip integrated linear-optical networks

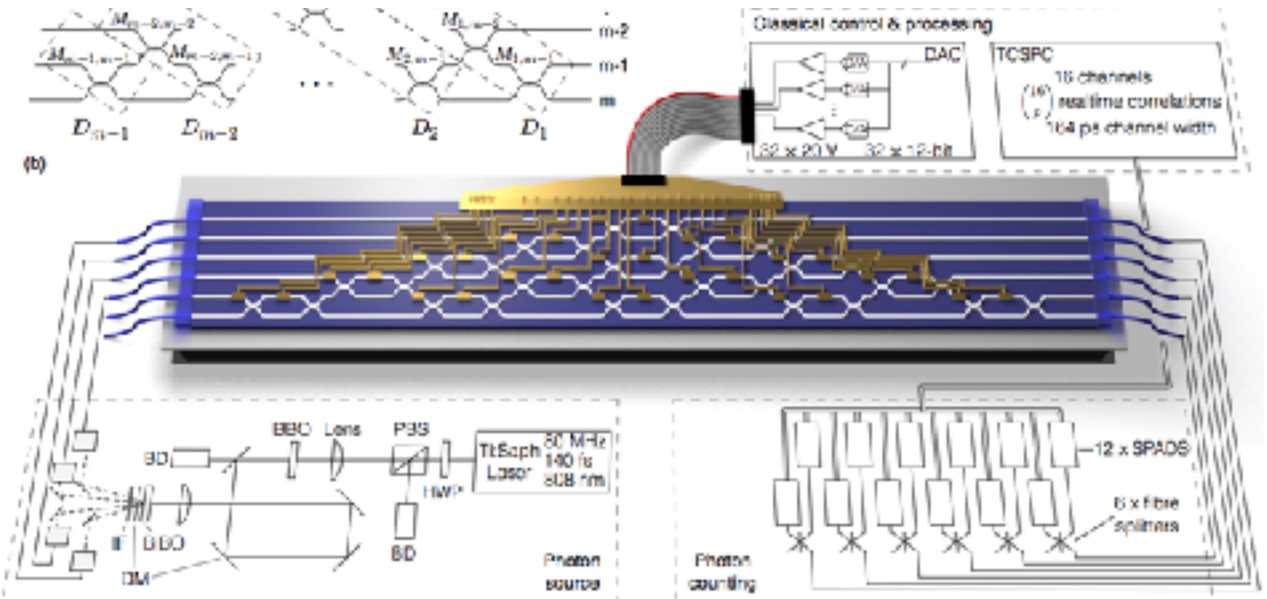

O’Brien; Walmsley; Walther; Sciarrino; White; etc.

Super-conducting qubit circuits

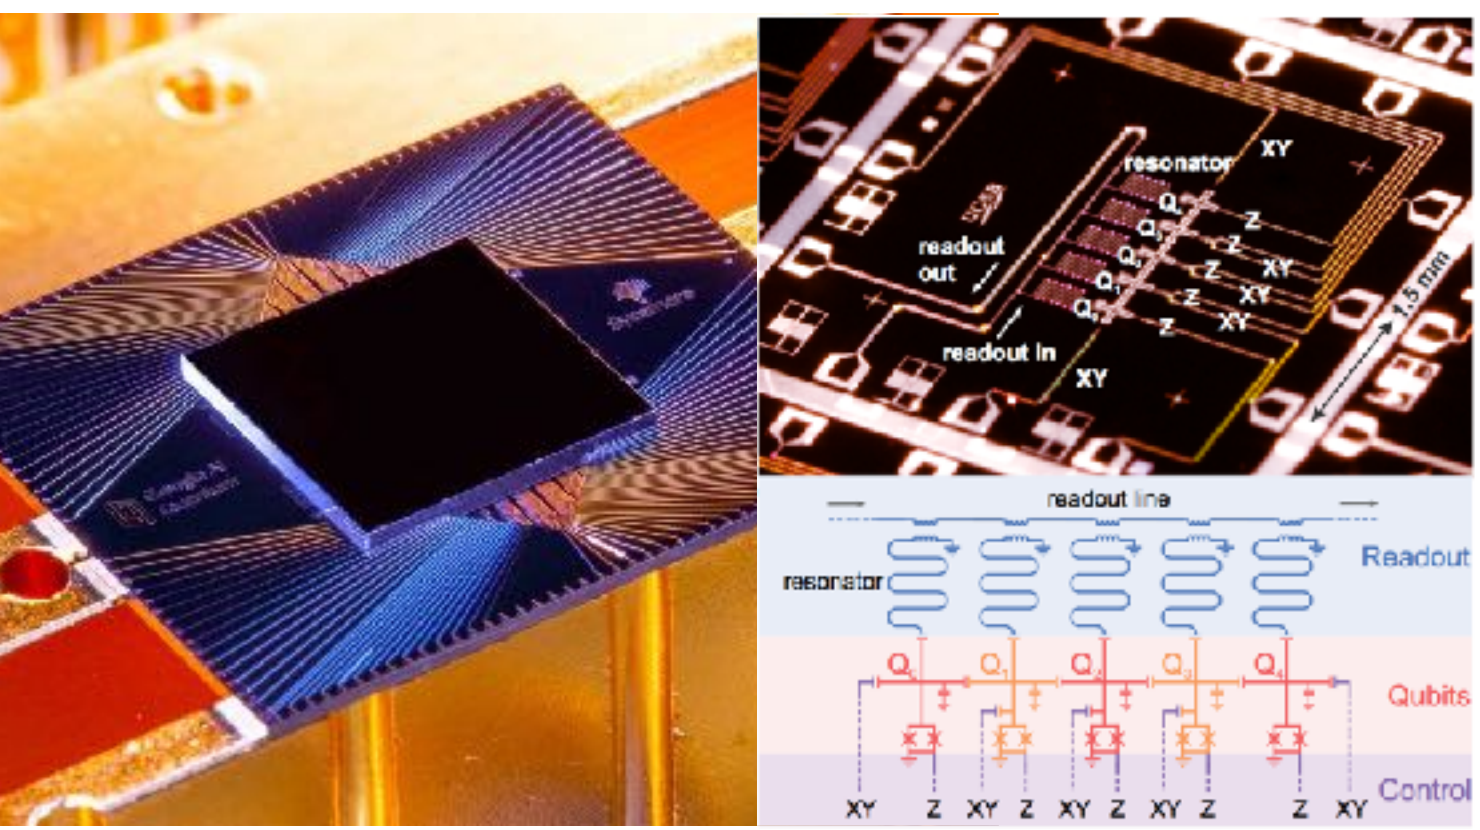

Martinis; Nori; Schoelkopf; Houk; Türeci; Blais; Wallraf; etc.

Cold atoms in optical lattices

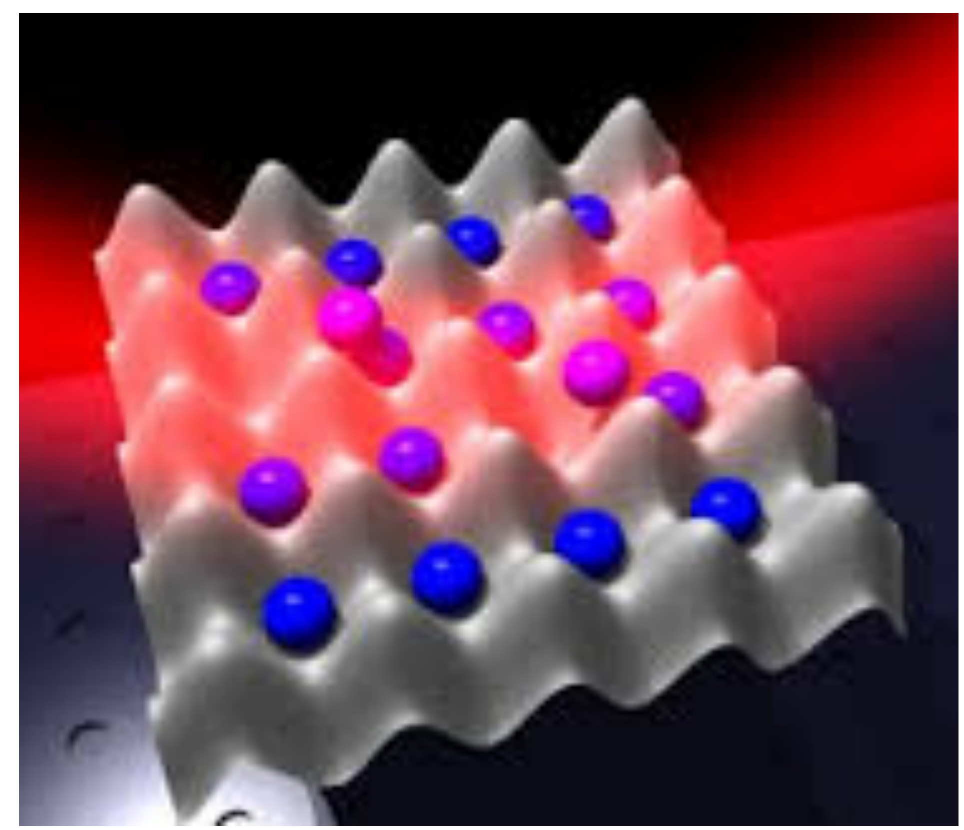

Bloch; Dalibard; Hänsch; Eisslinger; Greiner; Phillips; Porto; etc.

… but how do we trust the quantum devices we build?

State tomography is the gold standard for characterising quantum systems

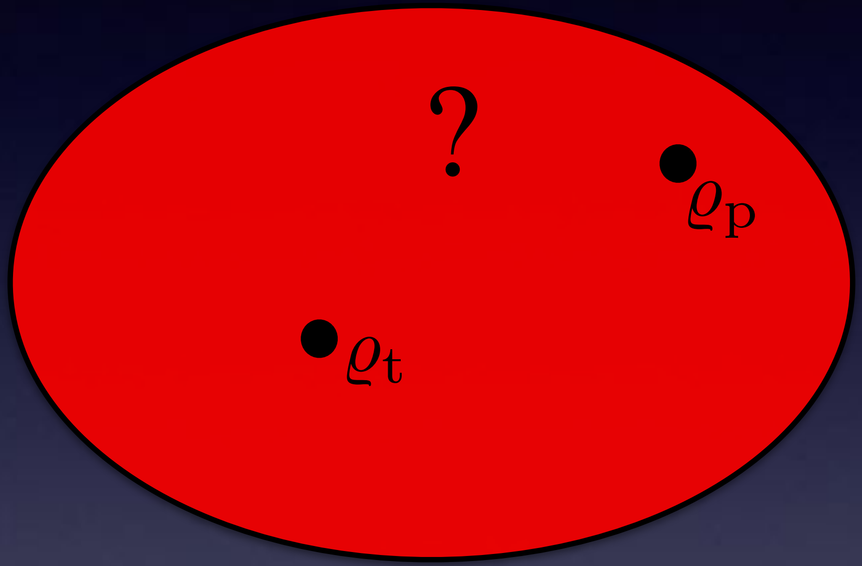

Where exactly is ϱ_{p}*?*

Hilbert space is a big place, but its physical corner is much smaller!

Physical systems occupy only a tiny (exponentially small) sector

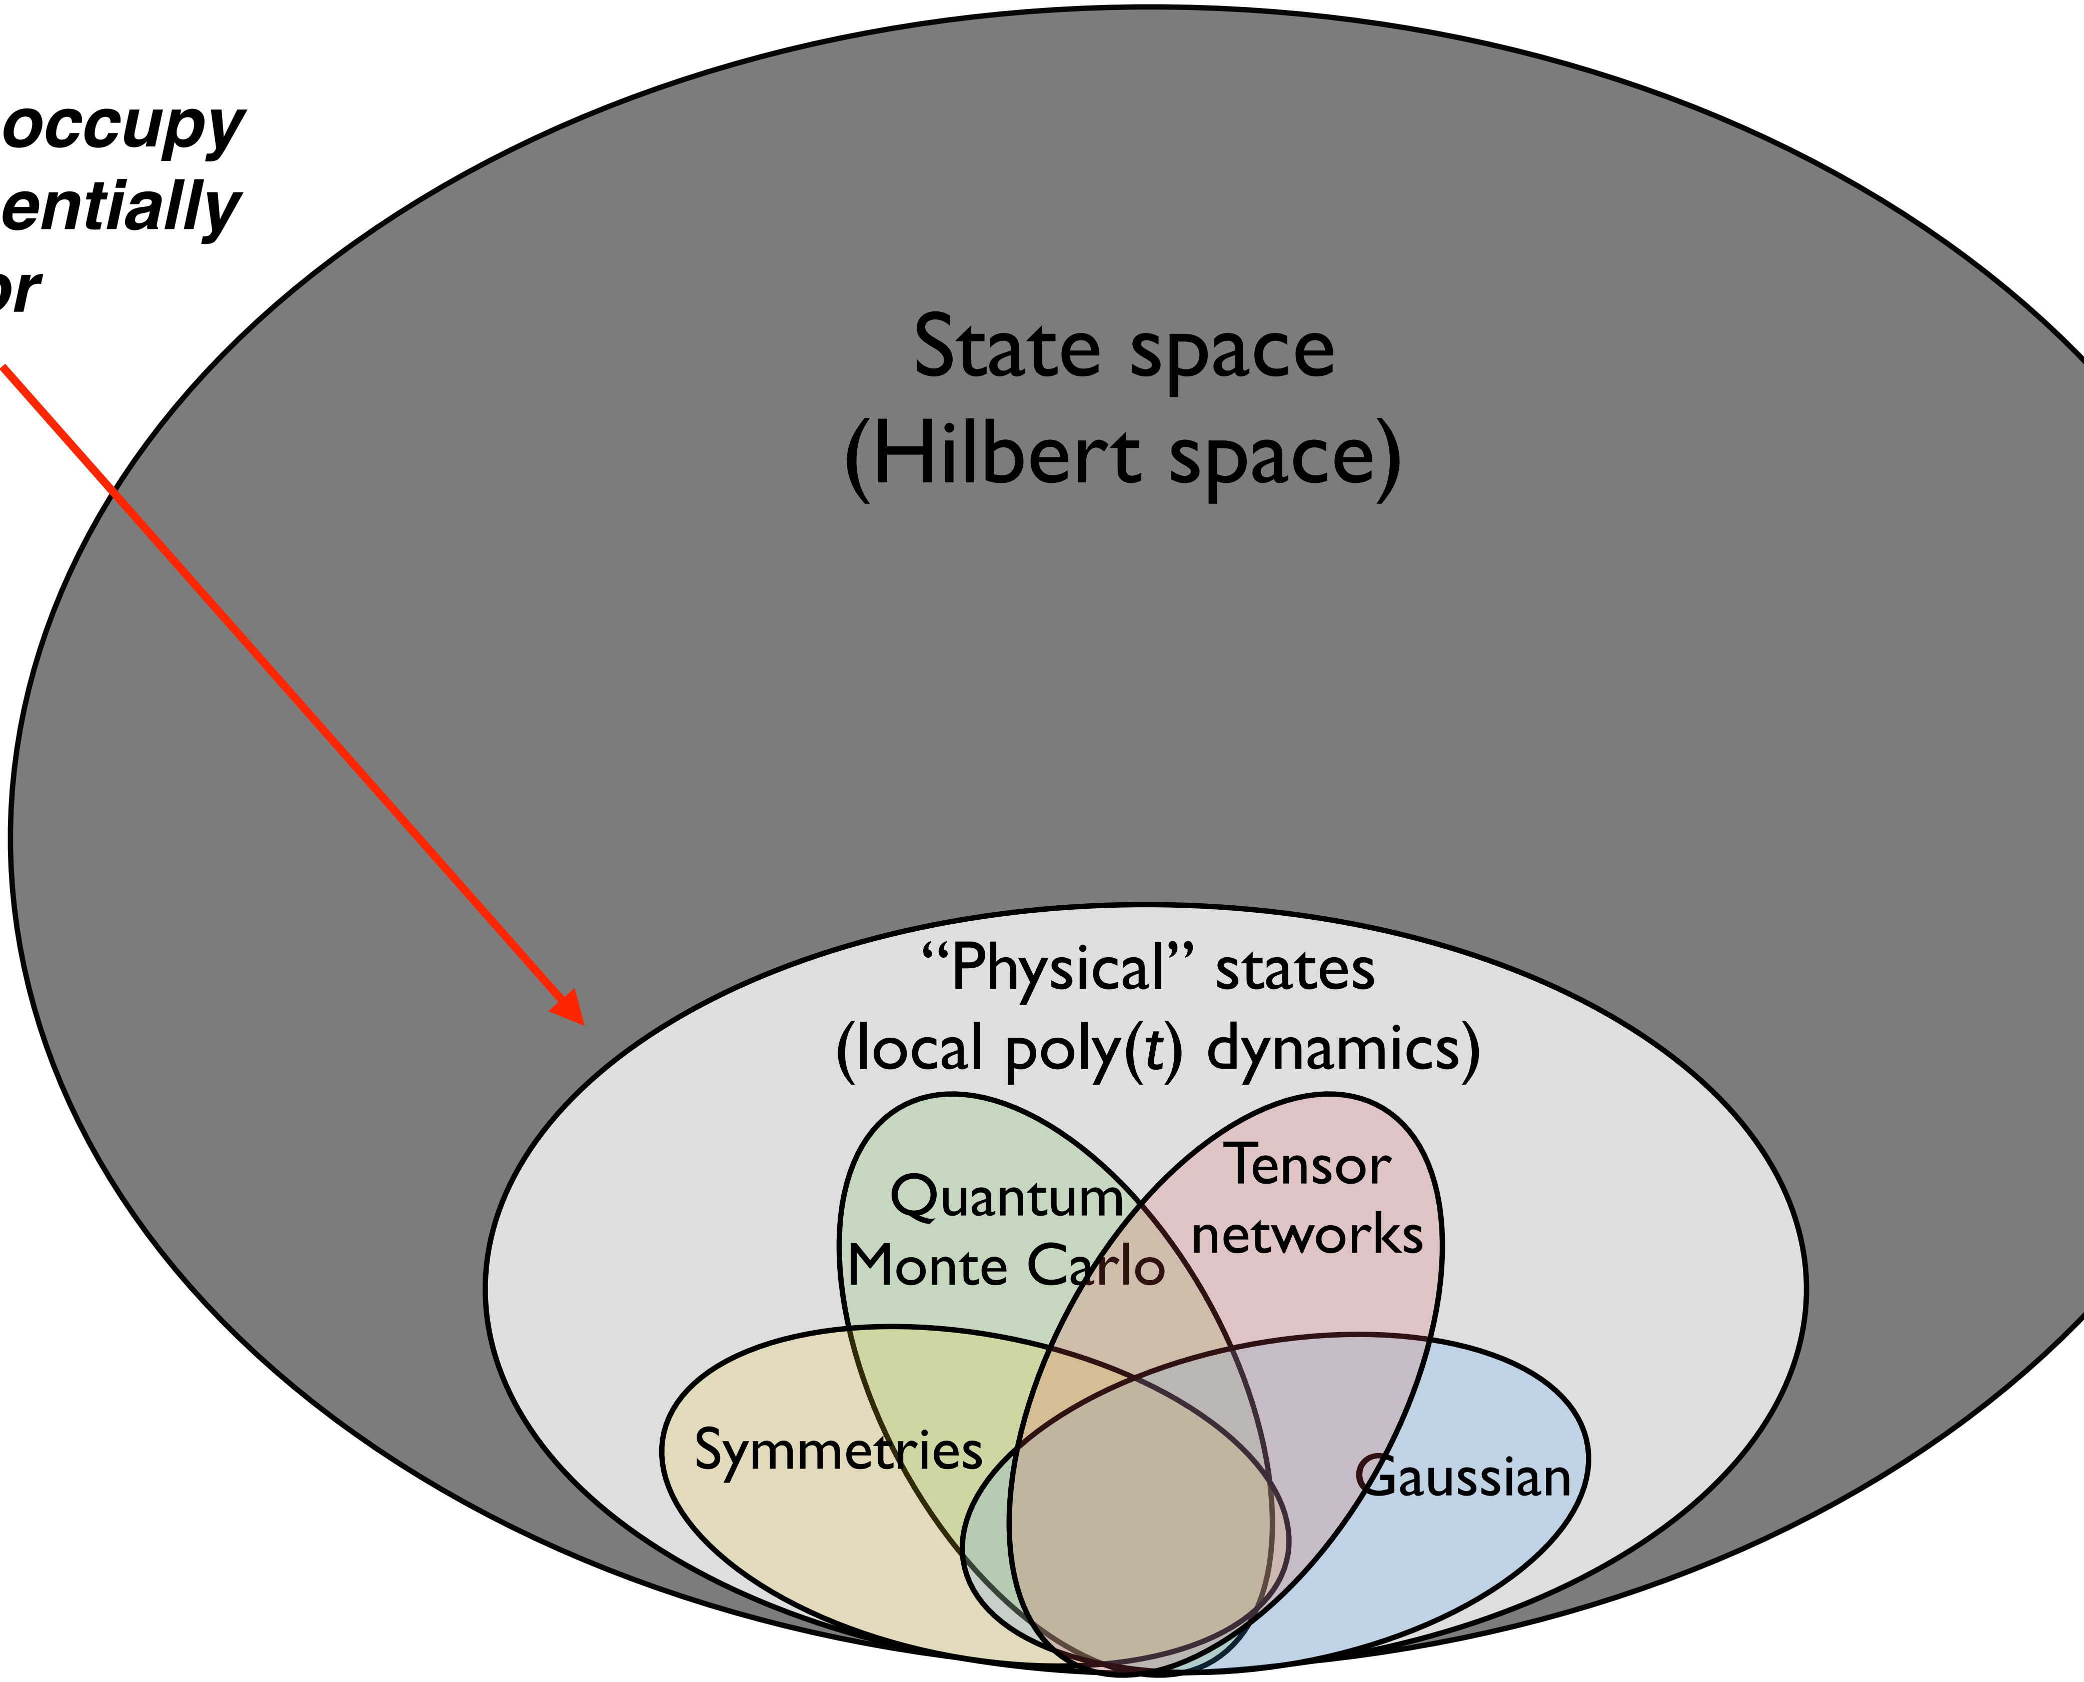

D. Poulin, A. Qarry, R. D. Somma, and F. Verstraete, PRL **106**, 170501 (2011).

Outline of the talk (Part I):

"Curse of dimensionality" also battled by the machine learning community

↓

Idea: quantum state and process tomography with classical neural nets

Part II: Classical simulations of quantum circuits.

Part III: Applications to near-term quantum computations.

Unsupervised learning and neural network (pure) states

- **Generative modelling**: Neural networks have proven successful at unsupervised learning unknown probability distributions from samples.

Hinton & Salakhutdinov, **Reducing the Dimensionality of Data with Neural Networks**, Science **313**, 504 (2006).

- **Neural-network states**: parametrize probability amplitudes with a neural-network Ansatz

G. Carleo and M. Troyer, Science **355**, 602 (2017);
Chen & Das Sarma, Phys. Rev. X (2017);
X. Gao and L.-M. Duan, Nat. Communs. **8**, 662 (2017);
I. Glasser et al., Phys. Rev. X **8**, 011006 (2018);
II. J. Chen et al., Phys. Rev. B **97**, 085104 (2018).

Idea: Born's probabilities as a neural net!

nature machine intelligence

ARTICLES

https://doi.org/10.1038/s42256-019-0028-1

Reconstructing quantum states with generative models

Juan Carrasquilla[1*], Giacomo Torlai[2,3,4], Roger G. Melko[2,3] and Leandro Aolita[5,6]

A major bottleneck in the development of scalable many-body quantum technologies is the difficulty in benchmarking state preparations, which suffer from an exponential 'curse of dimensionality' inherent to the classical description of quantum states. We present an experimentally friendly method for density matrix reconstruction based on neural network generative models. The learning procedure comes with a built-in approximate certificate of the reconstruction and makes no assumptions about the purity of the state under scrutiny. It can efficiently handle a broad class of complex systems including prototypical states in quantum information, as well as ground states of local spin models common to condensed matter physics. The key insight is to reduce state tomography to an unsupervised learning problem of the statistics of an informationally complete quantum measurement. This constitutes a modern machine learning approach to the validation of complex quantum devices, which may in addition prove relevant as a neural-network ansatz over mixed states suitable for variational optimization.

Reliable neural-network state tomography (artistic view)

- Parametrize ~~probability amplitudes~~ **measurement probabilities** with a **NN Ansatz**!
- Use **informationally-complete** (generalized) **measurements!**

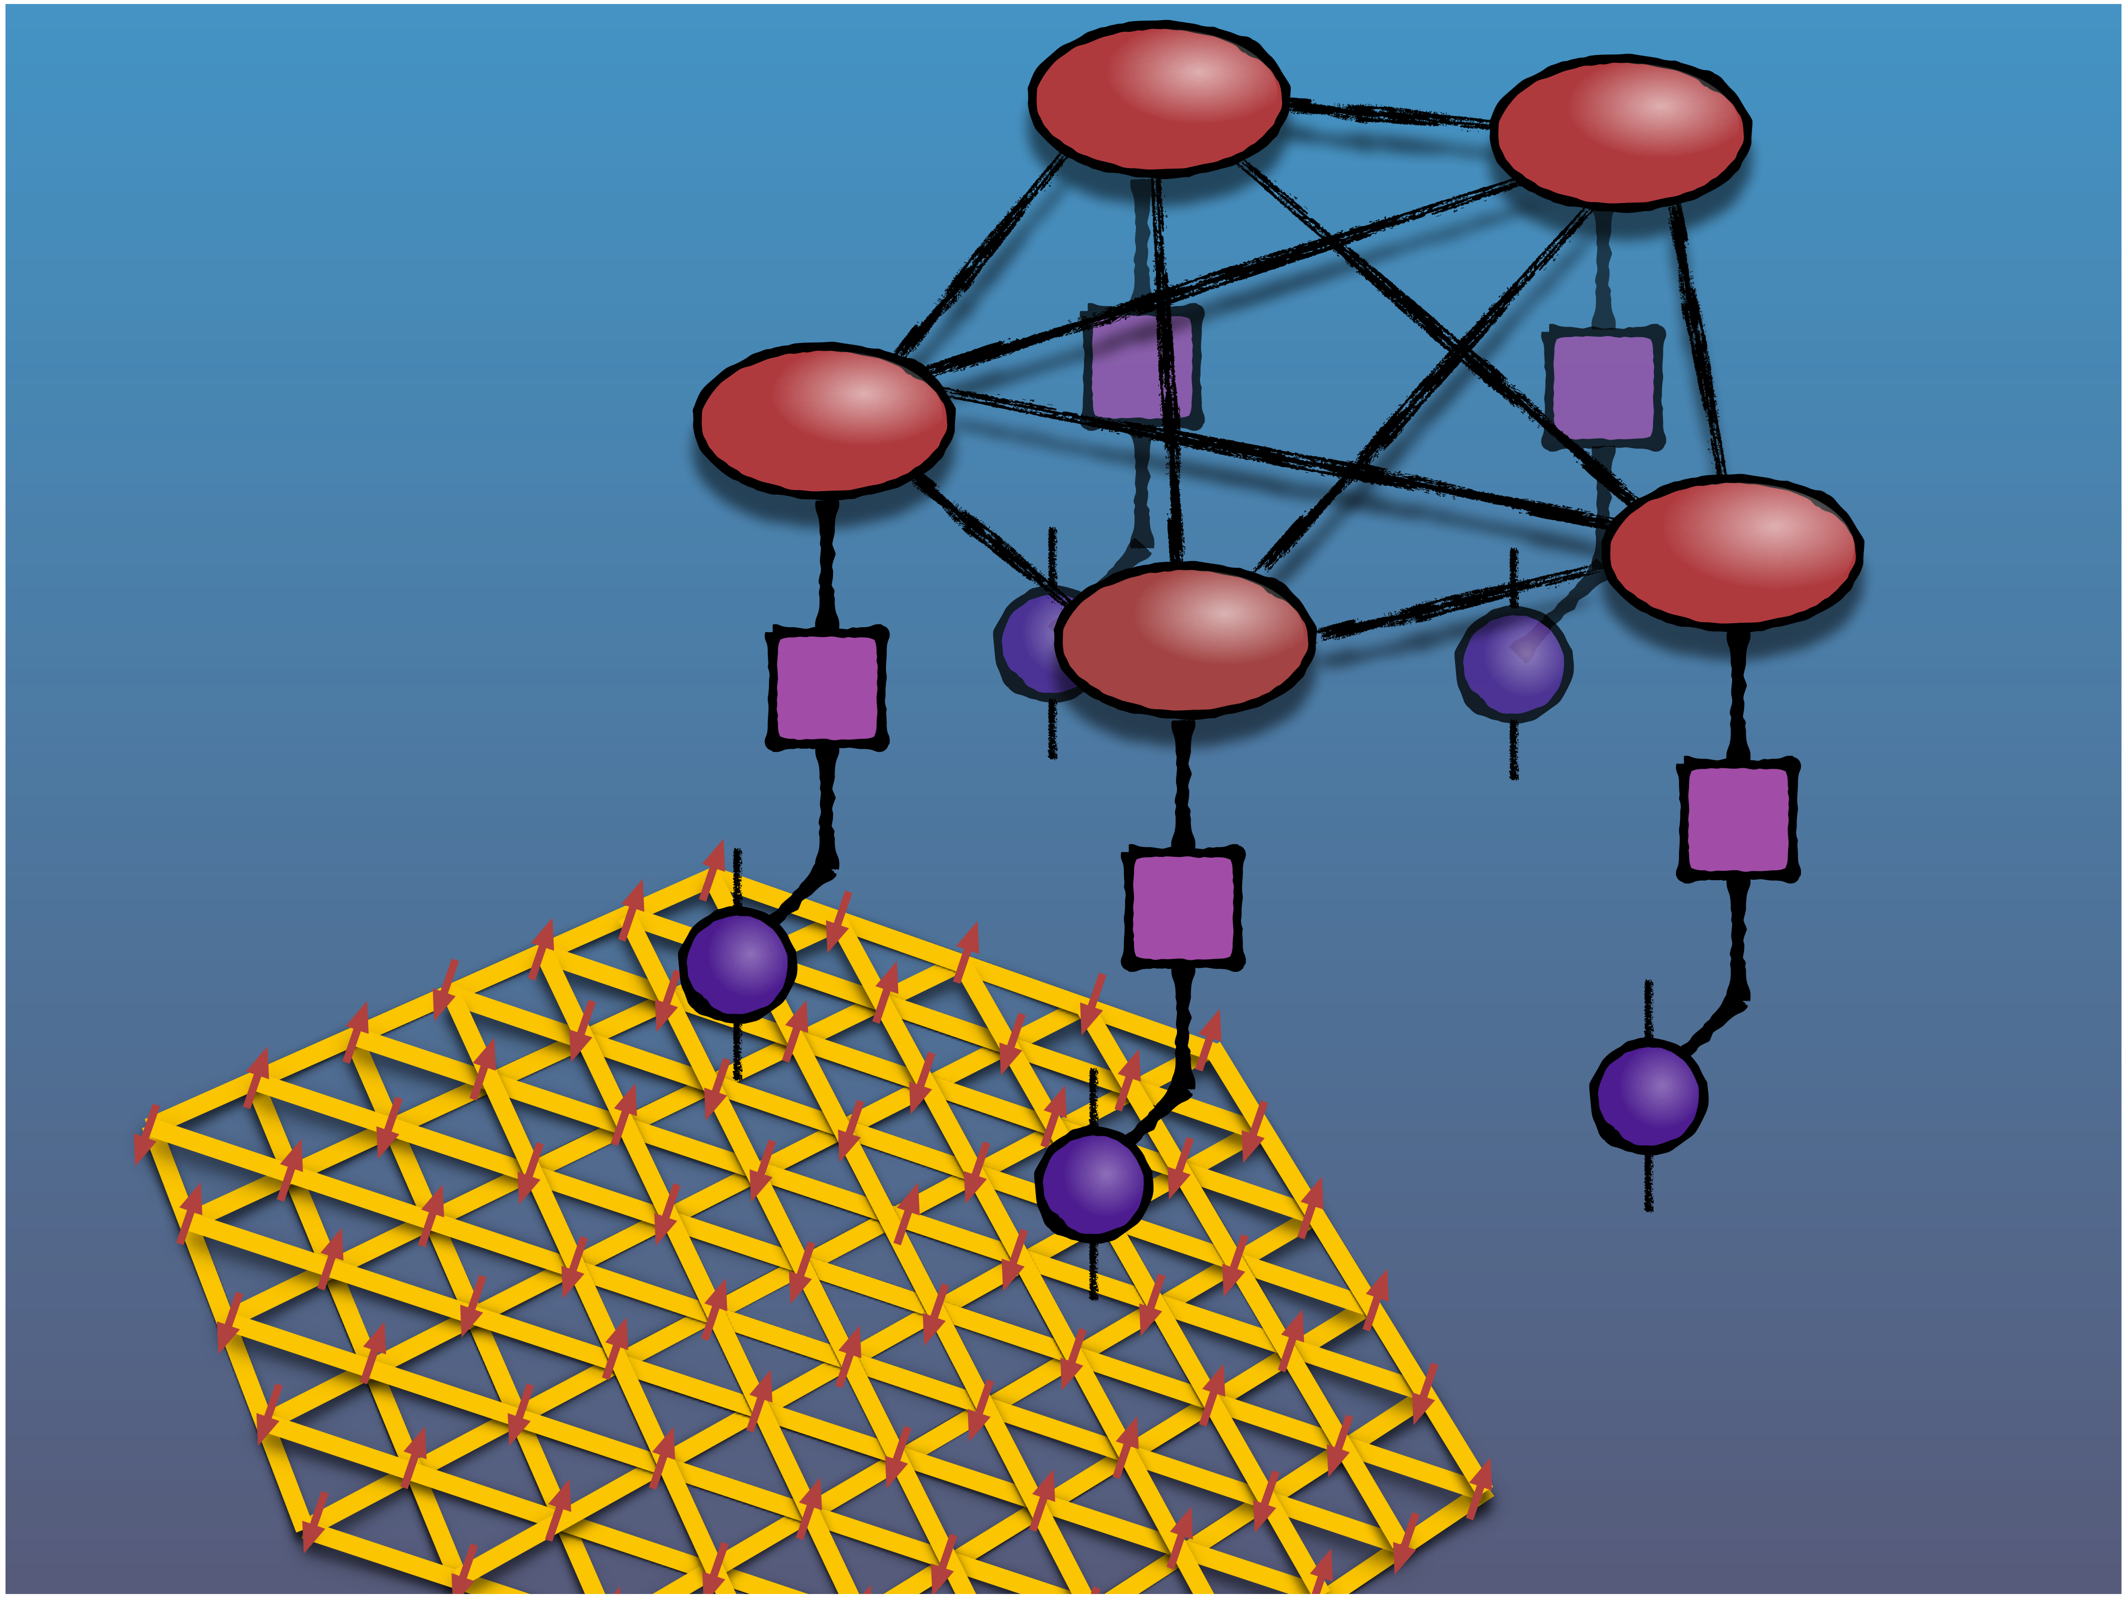

Reliable neural-network state tomography (actual representation)

$$\varrho_{\text{model}} = \left(\boldsymbol{P}_{\text{model}}\,\boldsymbol{T}^{-1}\right)^{t}\boldsymbol{M}$$

Highlights:

- Generic (**mixed) states encoded in a probability distribution**.
- **No density matrix stored**.
- **Monte-Carlo estimation of observables**: **Sampling + efficient tensor contractions**!

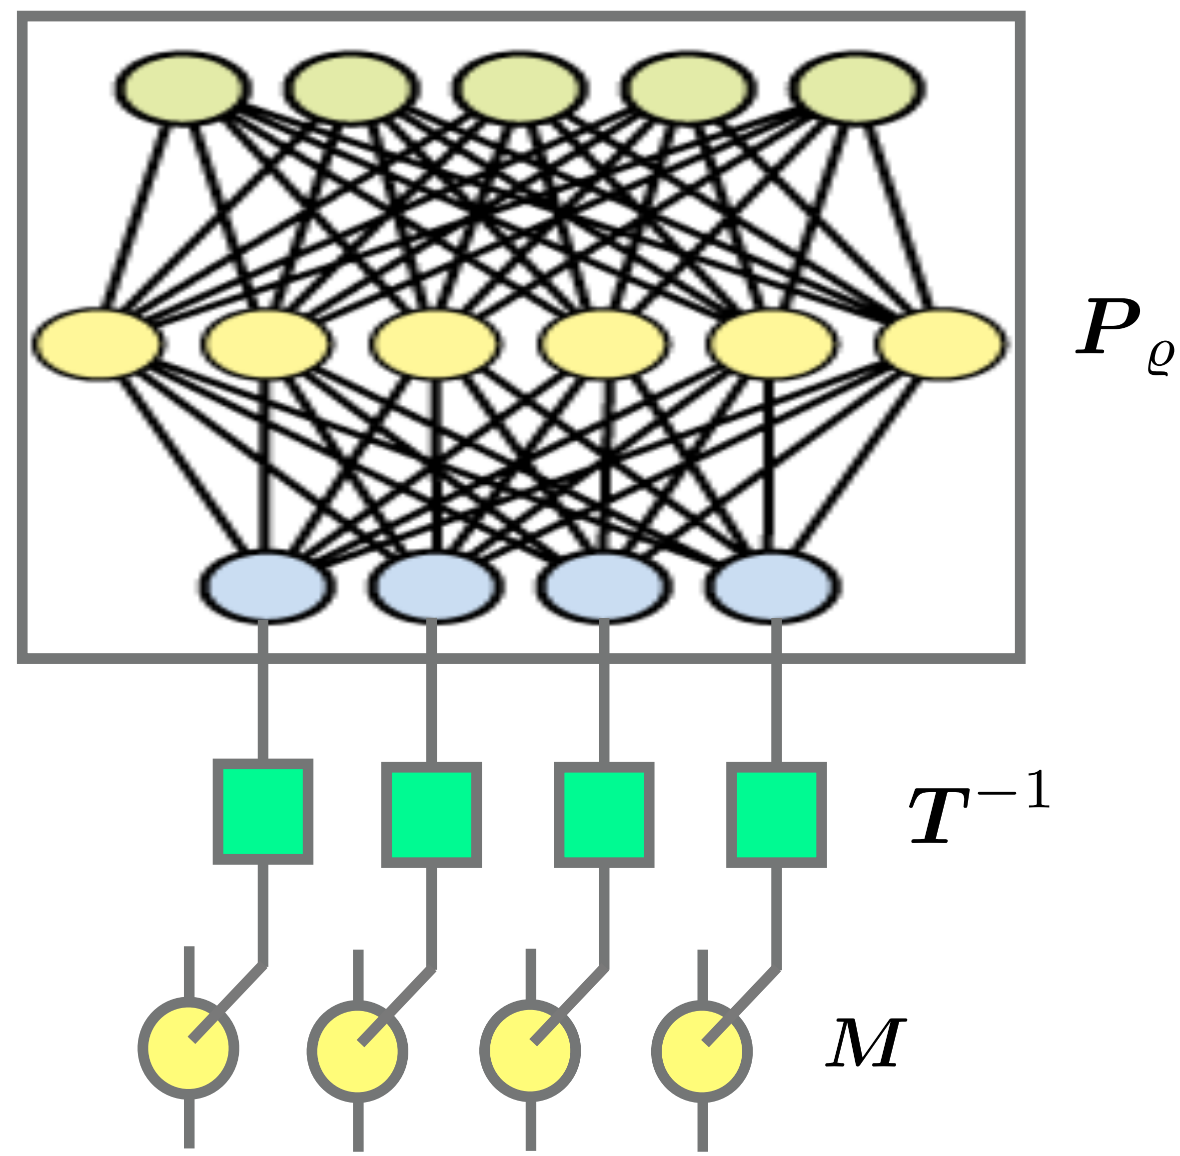

J. Carrasquilla, G. Torlai, R. Melko, and L. Aolita, **Reconstructing quantum states with generative models**, Nature Machine Intelligence **1**, 155 (2019).

Numerical experiments

AFM transverse-field Heisenberg-model ground state on a 2D lattice

8 x 8 spin-1/2 lattice. Deep (3 layer) recurrent neural network (GRU). IC measurement: tetrahedron or Pauli-eigenstates:

$\langle \boldsymbol{\sigma}_1 . \boldsymbol{\sigma}_j \rangle$ *RNN*

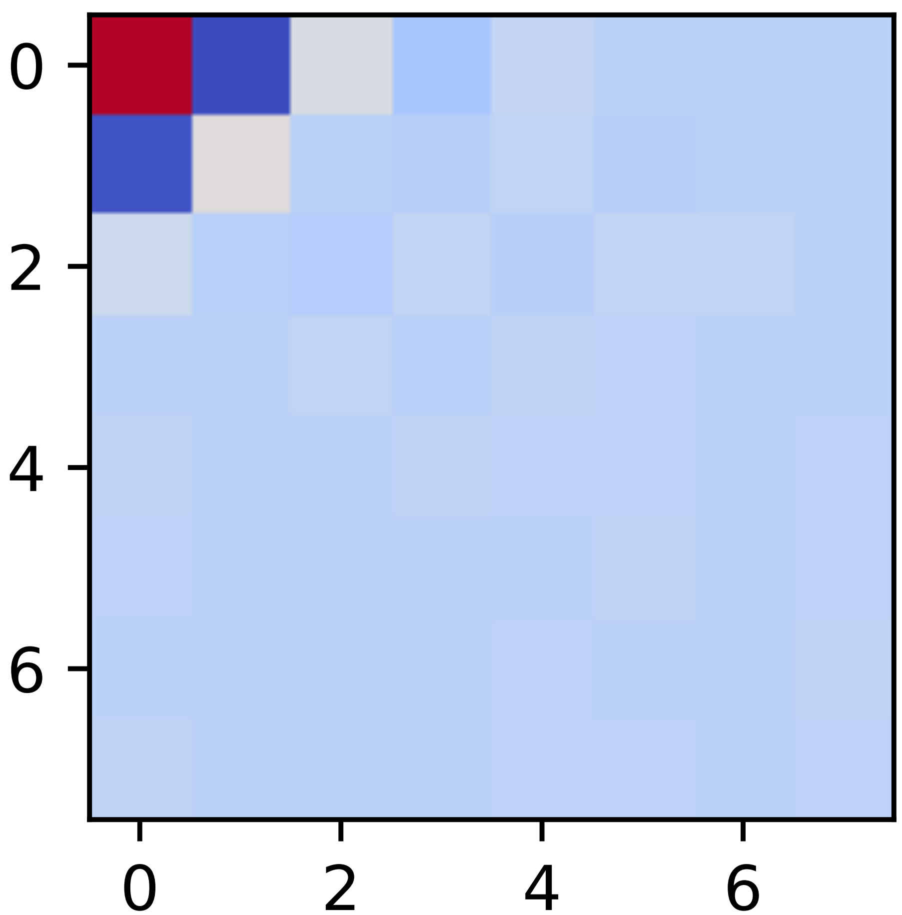

$\langle \boldsymbol{\sigma}_1 . \boldsymbol{\sigma}_j \rangle$ *2D-DMRG*

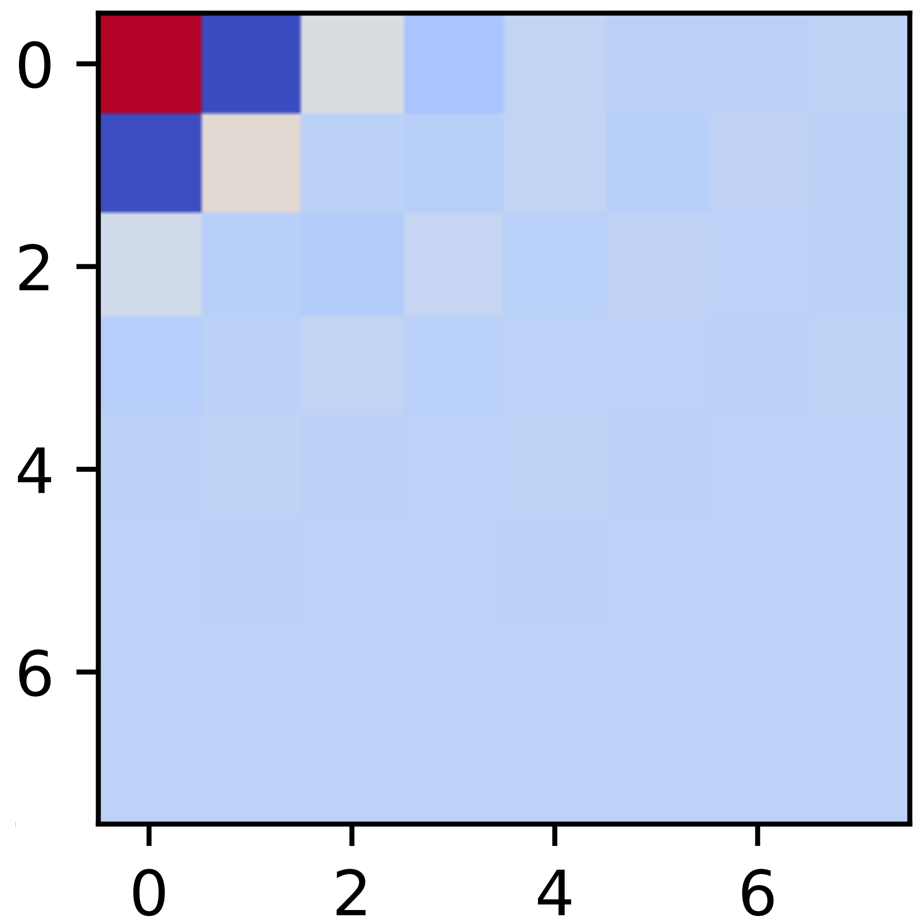

$$H = J \sum_{i,j} \boldsymbol{\sigma}_i . \boldsymbol{\sigma}_j$$

Heisenberg model

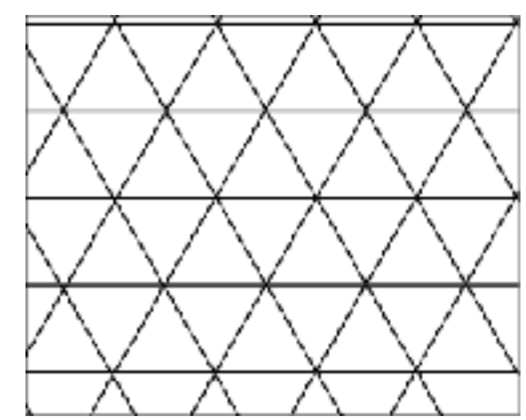

Triangular lattice used

- Model learns **(F_C = 0.998)** wave function with **complex sign structure!**
- **Out of reach for MPS tomography.**
- **Efficient estimation** of expectation values **of two-body observables**.

Quantum process tomography even more challenging

G. Torlai, C. J. Wood, A. Acharya, G. Carleo, J. Carrasquilla, and LA, **Quantum process tomography with unsupervised learning and tensor networks**, arXiv: 2006.02424.

Process reconstruction of random quantum circuits

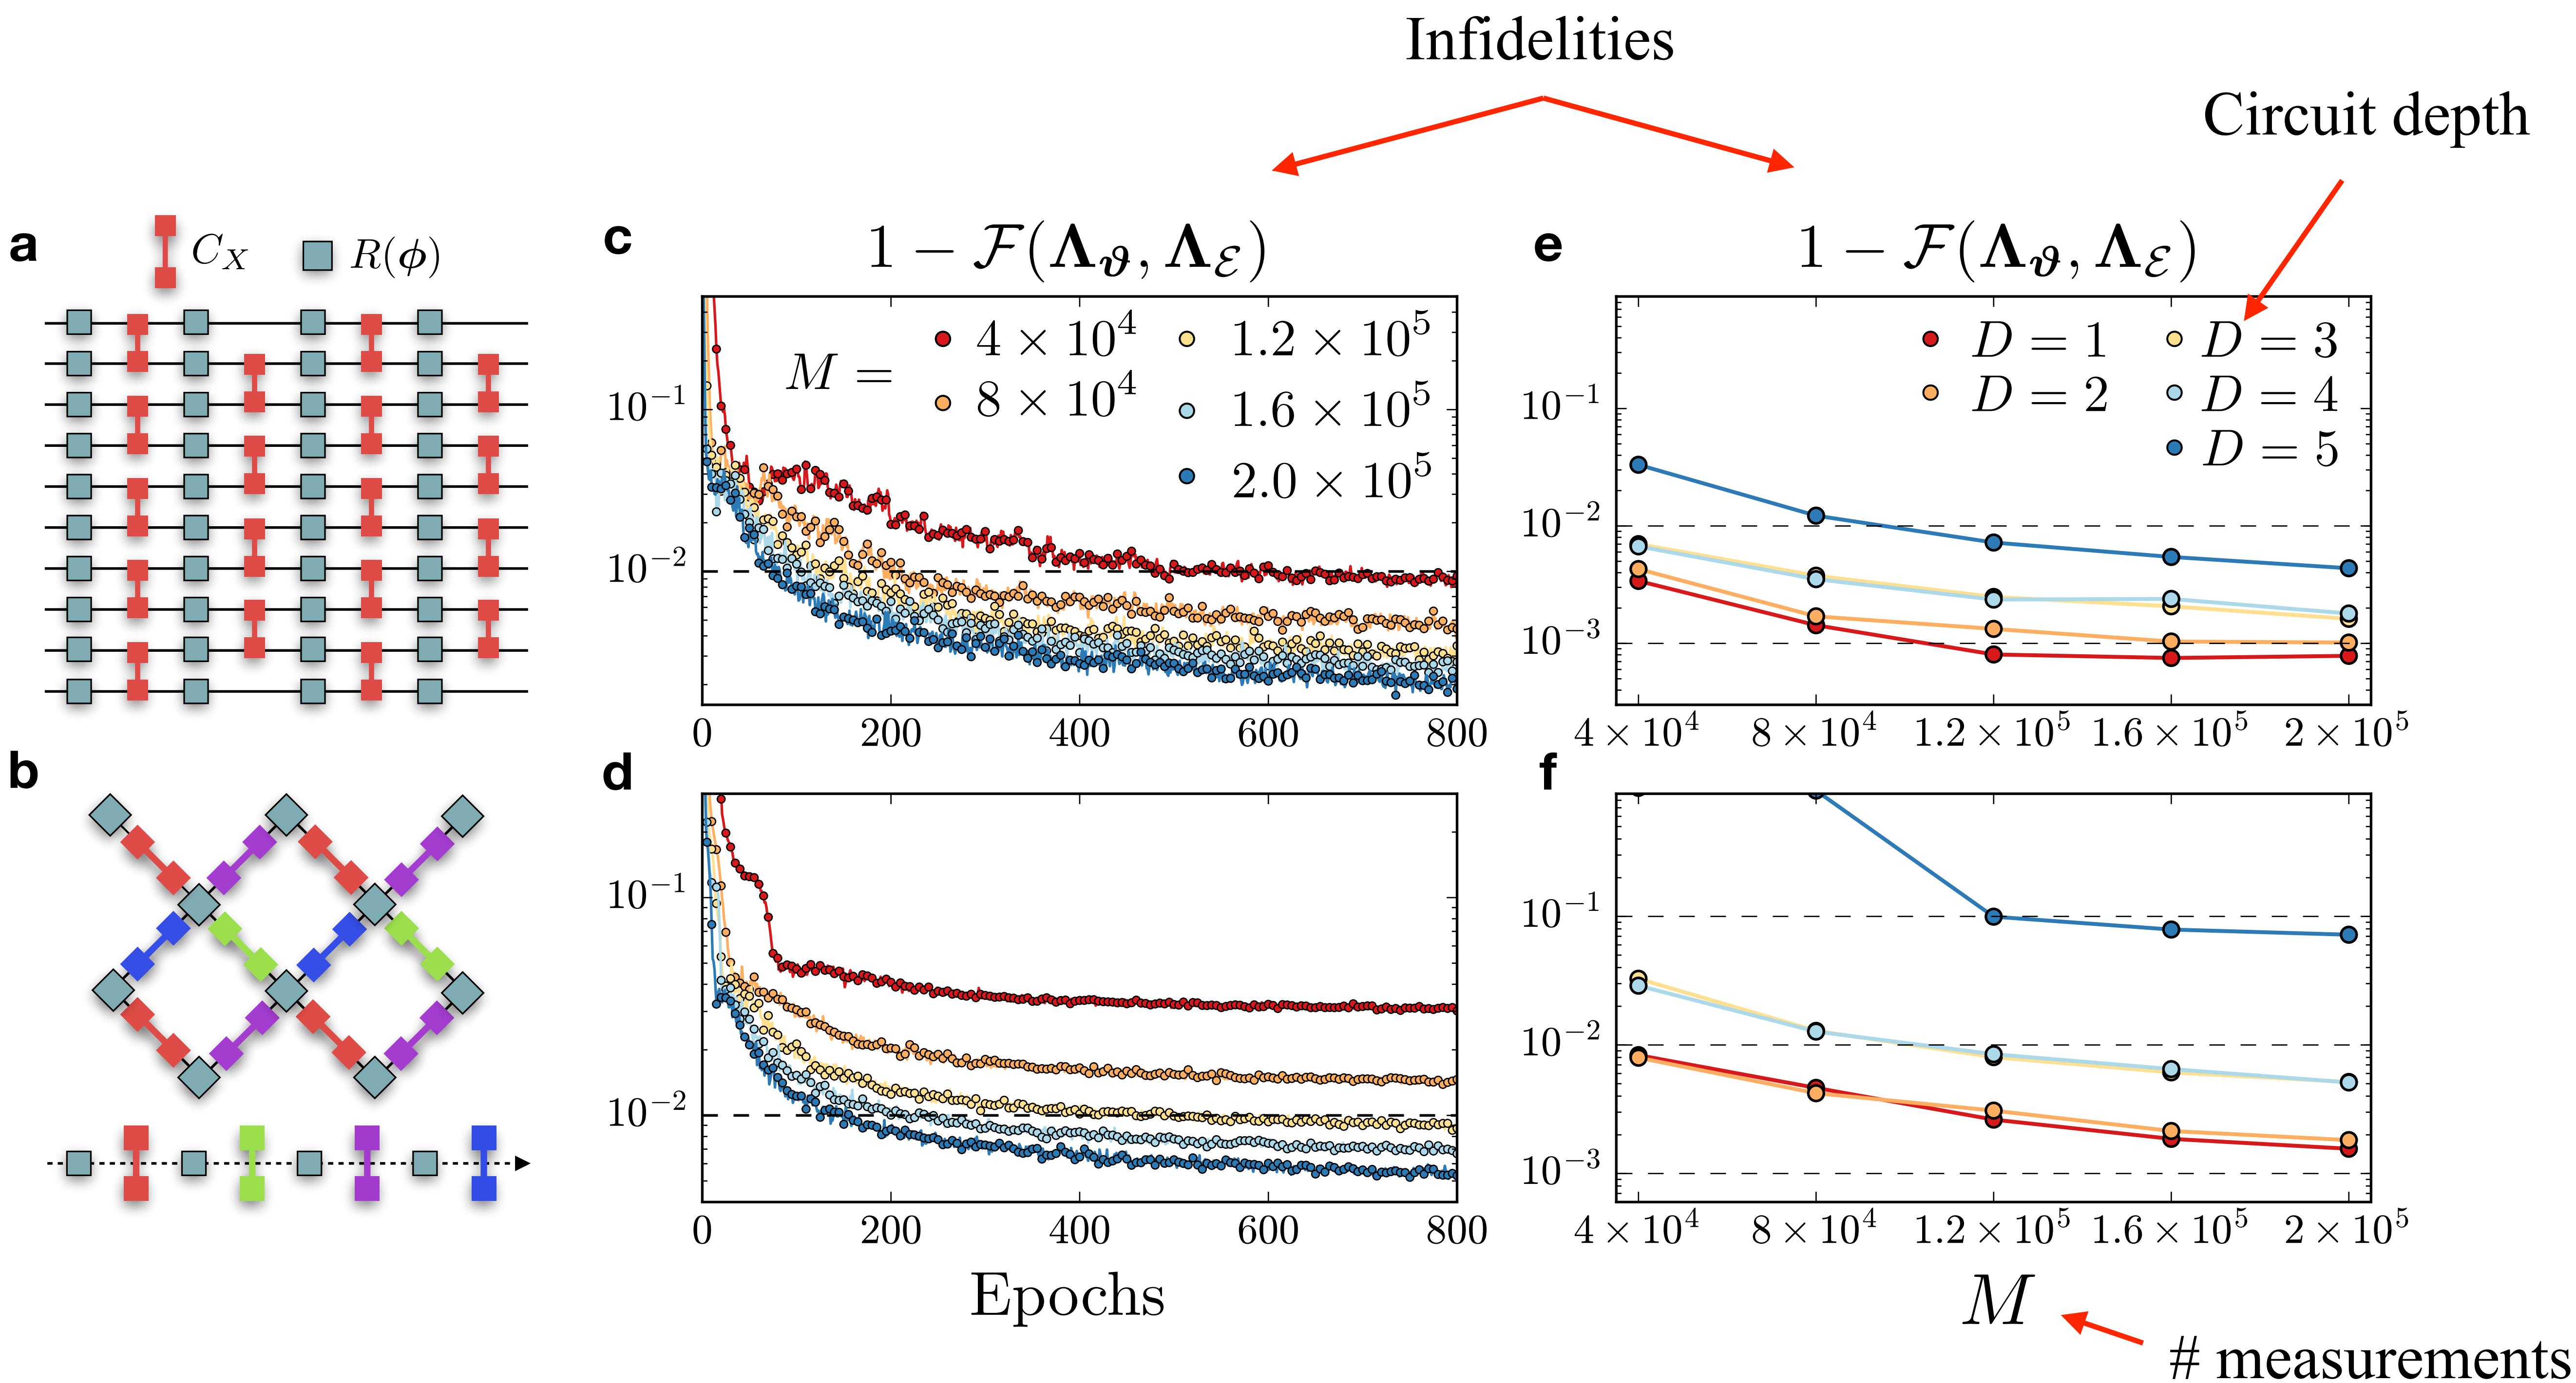

- Random quantum circuits are **computationally hard**.
- **Tensor-network generative model** learns the processes for ***N*=10 qubits and depth *D*=5.**
- **Previous reconstructions:** simple gates involving **only *N*=3 qubits.**
- **Full process tomography for *N*=10 would require ~10^{12} settings!**

Runtime scalings of quantum process learning

Reconstruction fidelity versus **classical processing time** for a **10-qubit depth-5** random quantum circuit (**on a laptop!**), data size M=100 k samples:

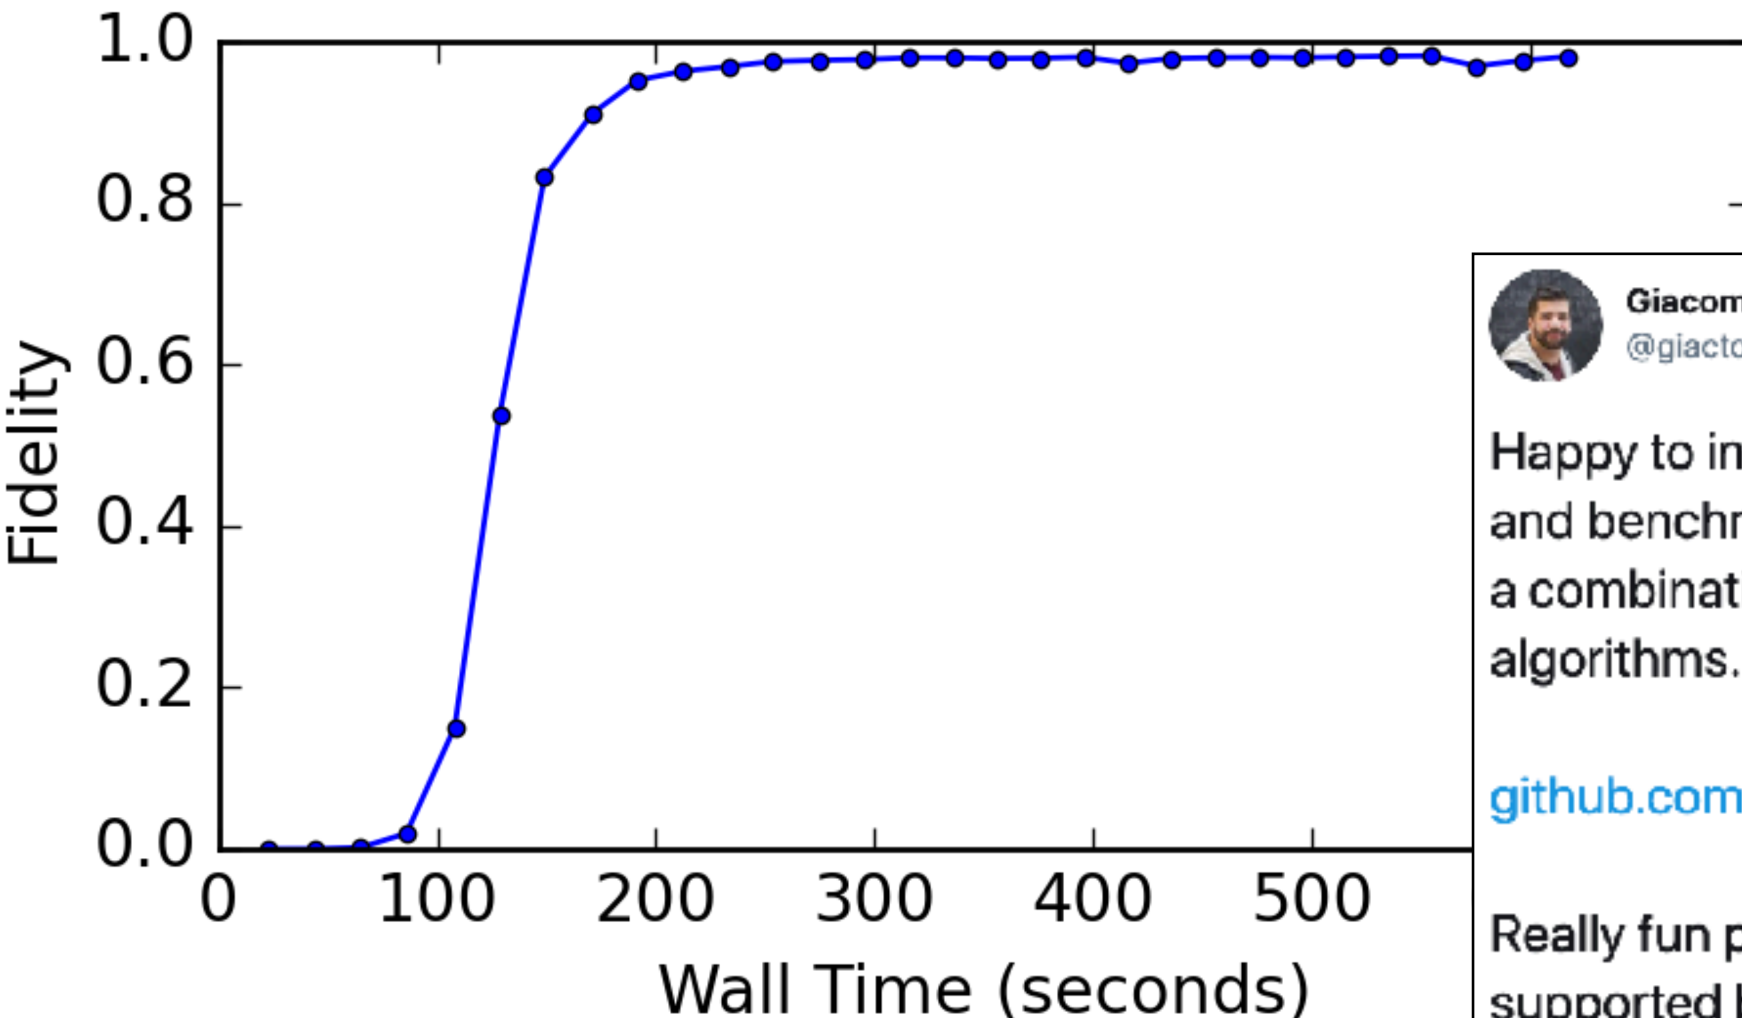

Part I´s Conclusions:

- Idea: neural Ansatz for the measurement statistics:

1. Record-breaking performances for state and process reconstructions.
2. Unsupervised learning + tensor networks + Monte-Carlo.
3. Efficient estimation of observables (state bypassed!)
4. Noise characterisations useful for error-correction schemes.

J. Carrasquilla, G. Torlai, R. Melko, and L. Aolita, **Reconstructing quantum states with generative models**, Nature Machine Intelligence **1**, 155 (2019).

G. Torlai, C. J. Wood, A. Acharya, G. Carleo, J. Carrasquilla, and LA, **Quantum process tomography with unsupervised learning and tensor networks**, arXiv: 2006.02424.

Part II: simulations of quantum circuits

A machine-learning amenable formulation of quantum mechanics

J. Carrasquilla, D. Luo, F. Pérez, A. Milstead, B. K. Clark, M. Volkovs, and LA,
Probabilistic Simulation of Quantum Circuits with the Transformer,
arXiv: 1912.11052.

A probabilistic formulation of quantum dynamics

Algorithm: Given a state's distribution, sample from the updated distribution after a local circuit gate. Then, learn the representation of the updated state from those samples.

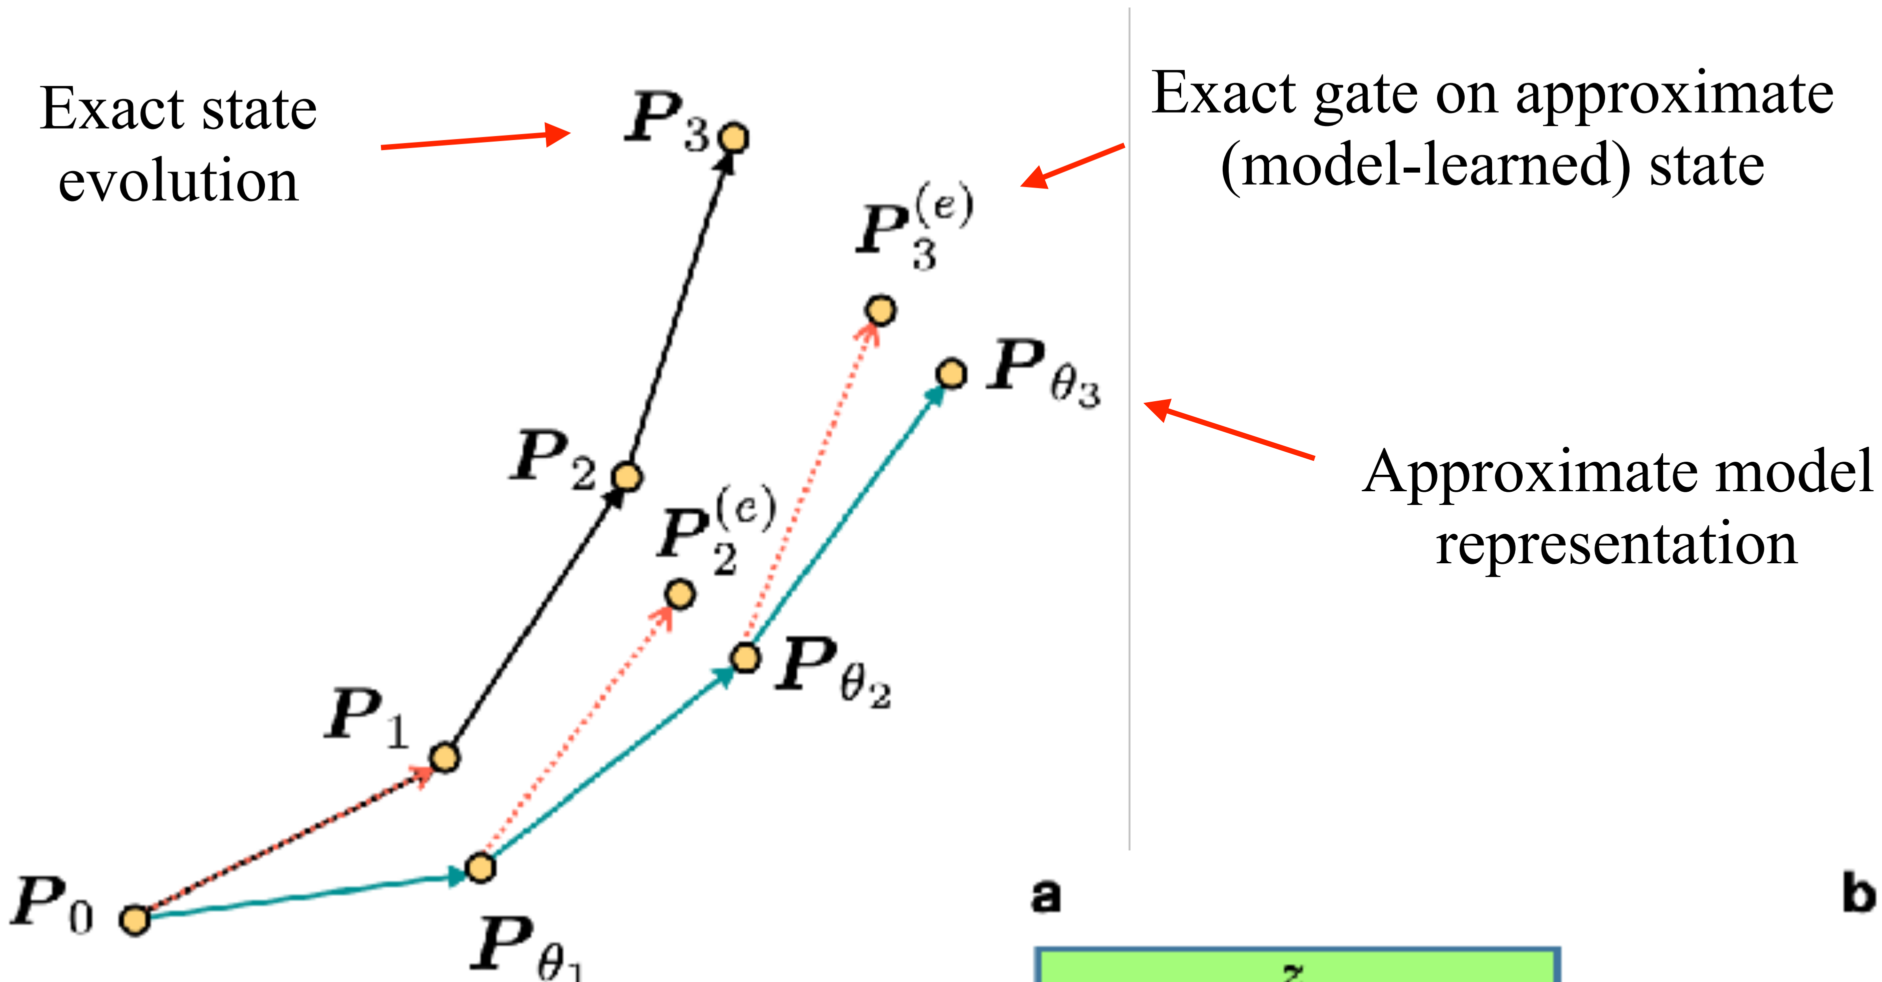

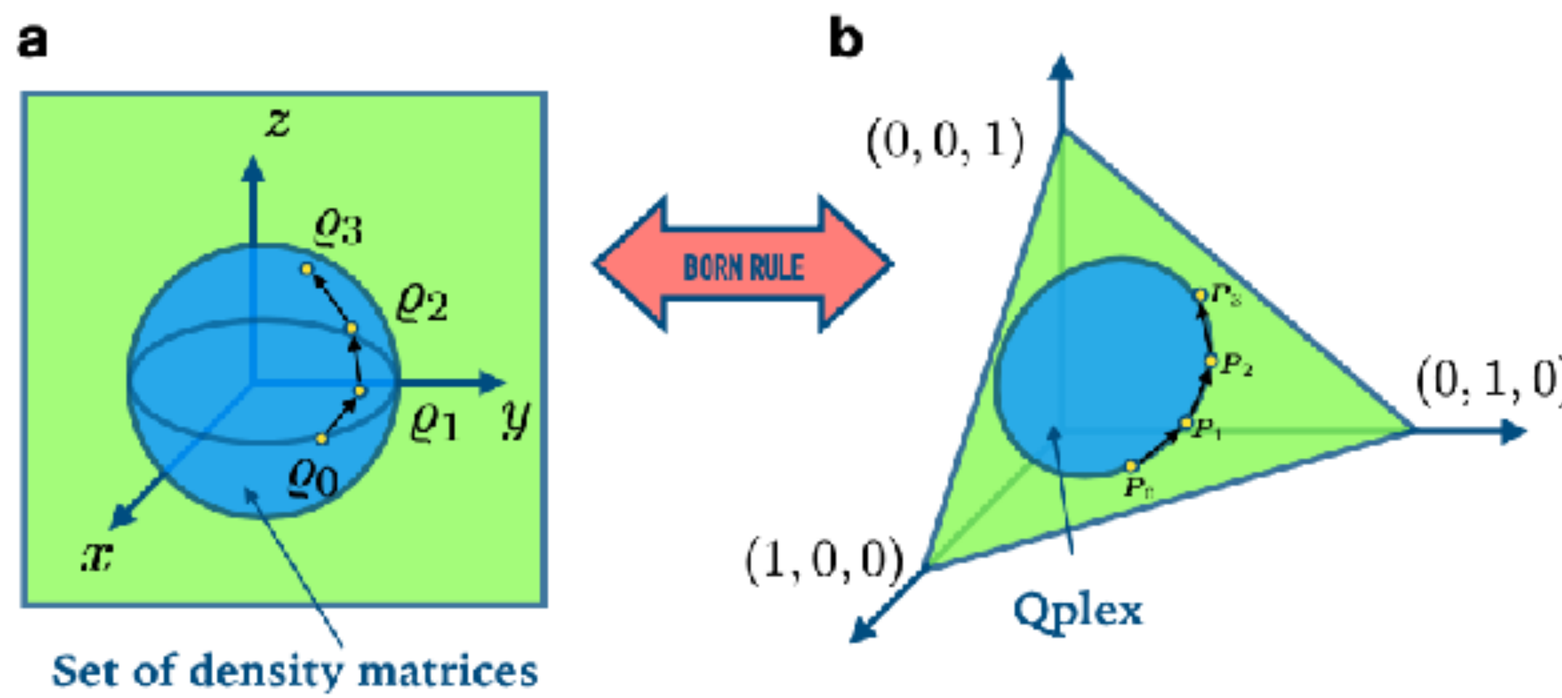

Classical simulation of quantum circuits with "the Transformer"

The Transformer: powerful generative model with a **tractable** probability **distribution** and that allows for **exact sampling**.

Fidelities for GHZ and graph state circuits

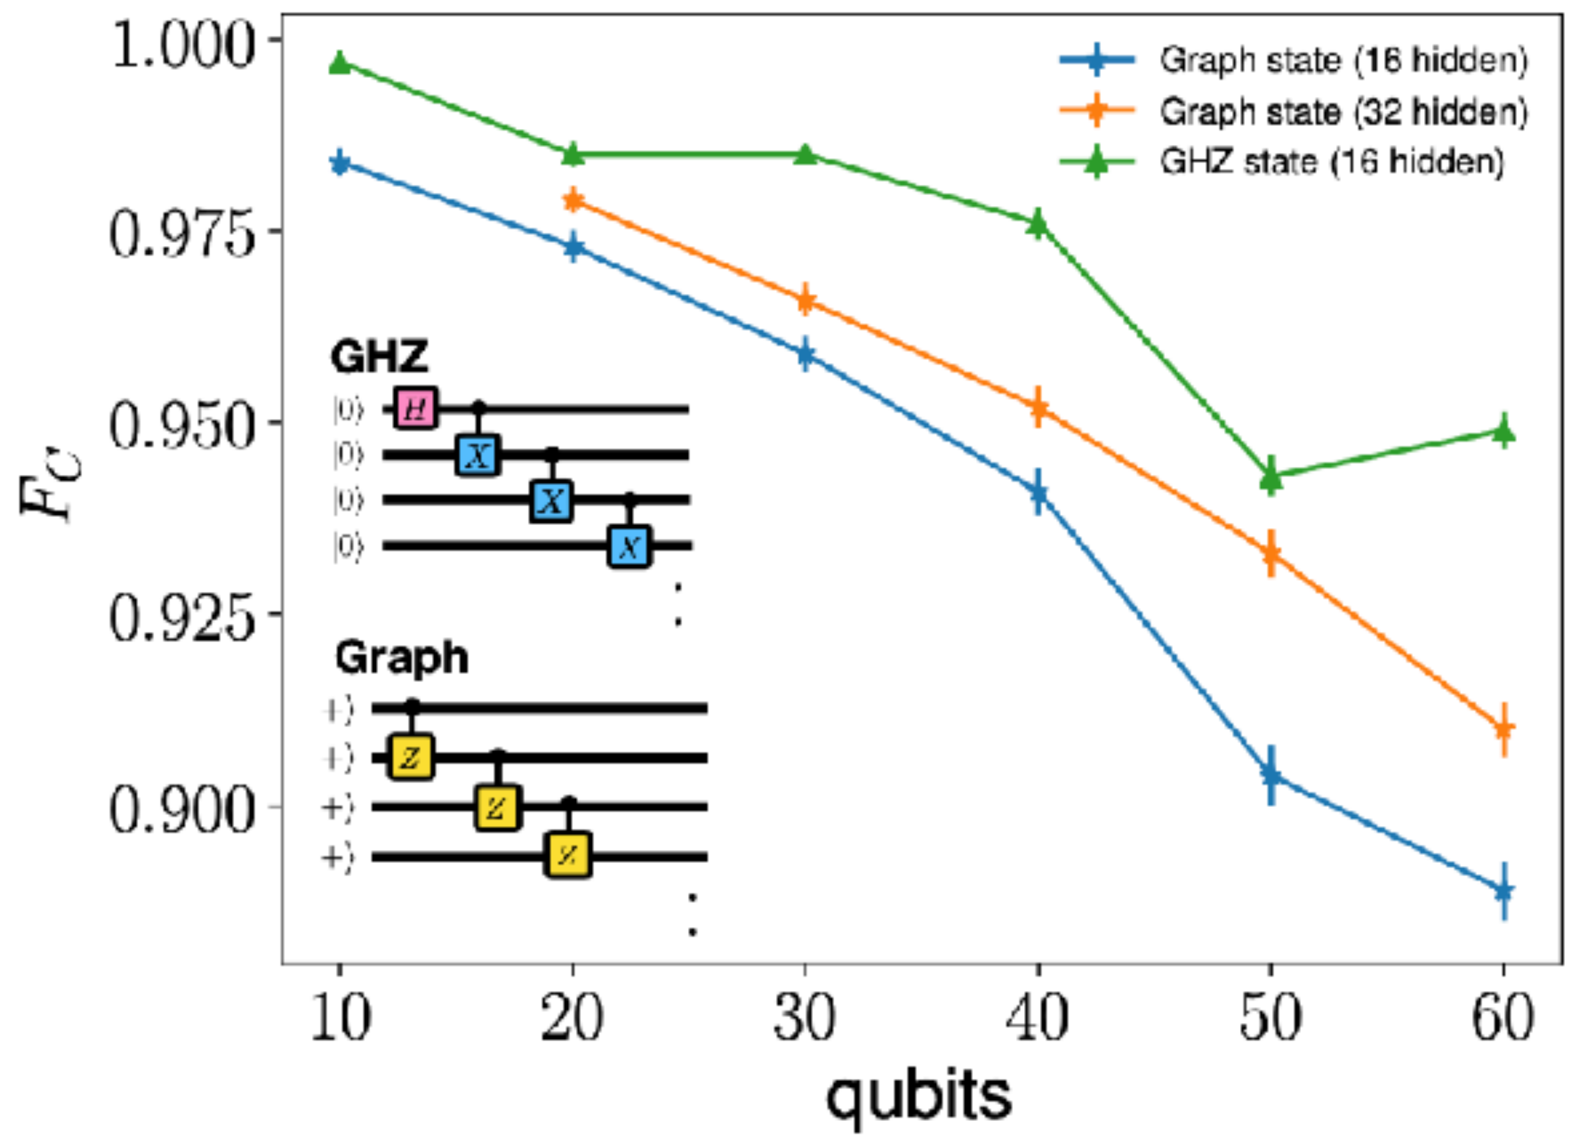

- **Exact sampling and tractable distributions** significantly simplifies computations (no Monte-Carlo Markov chain required!).
- Only a **proof-of-principle**, optimisations for considerable **efficiency increase possible.**


An alternative we did not anticipate

Quantum states encoded on neuromorphic circuits

S. Czischek, A. Baumbach, S. Billaudelle, B. Cramer, L. Kades, J. M. Pawlowski, M. Oberthaler, J. Schemmel, M. A. Petrovici, T. Gasenzer, and M. Gärttner, **Spiking neuromorphic chip learns entangled quantum states**, arXiv: 2008.01039.

For neural-network software, better use neuromorphic hardware!

Structural and dynamical properties of biological neuronal networks with the aim of inheriting the brain's functional performance and energy efficiency.

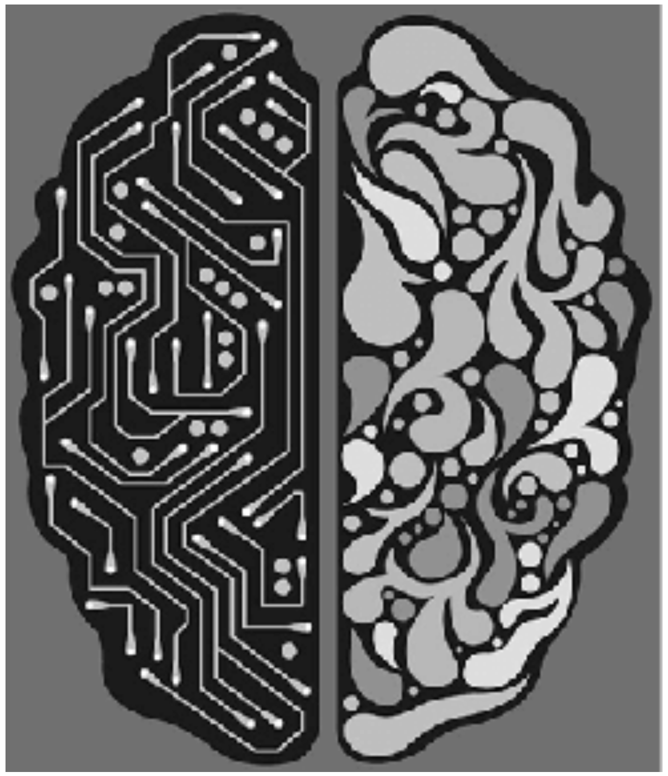

Neuromorphic device

von Neumann computer

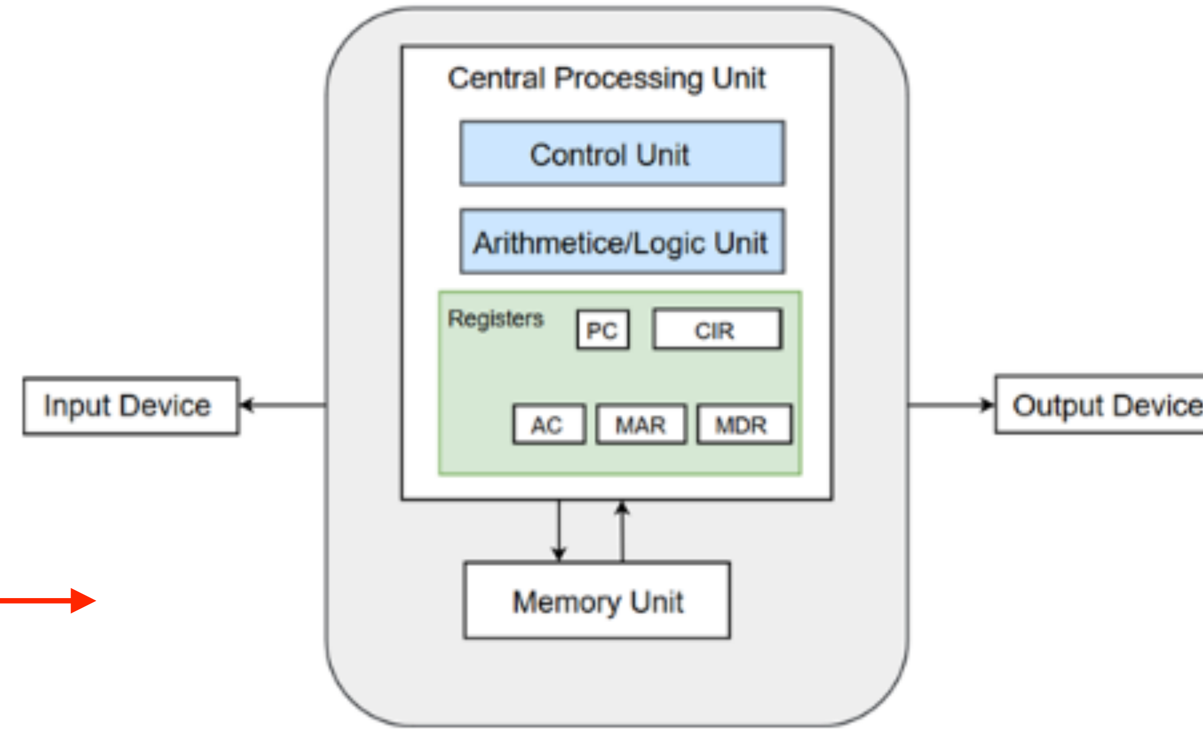

Brain-like architectures and neural-tailored devices will allow for smarter algorithms

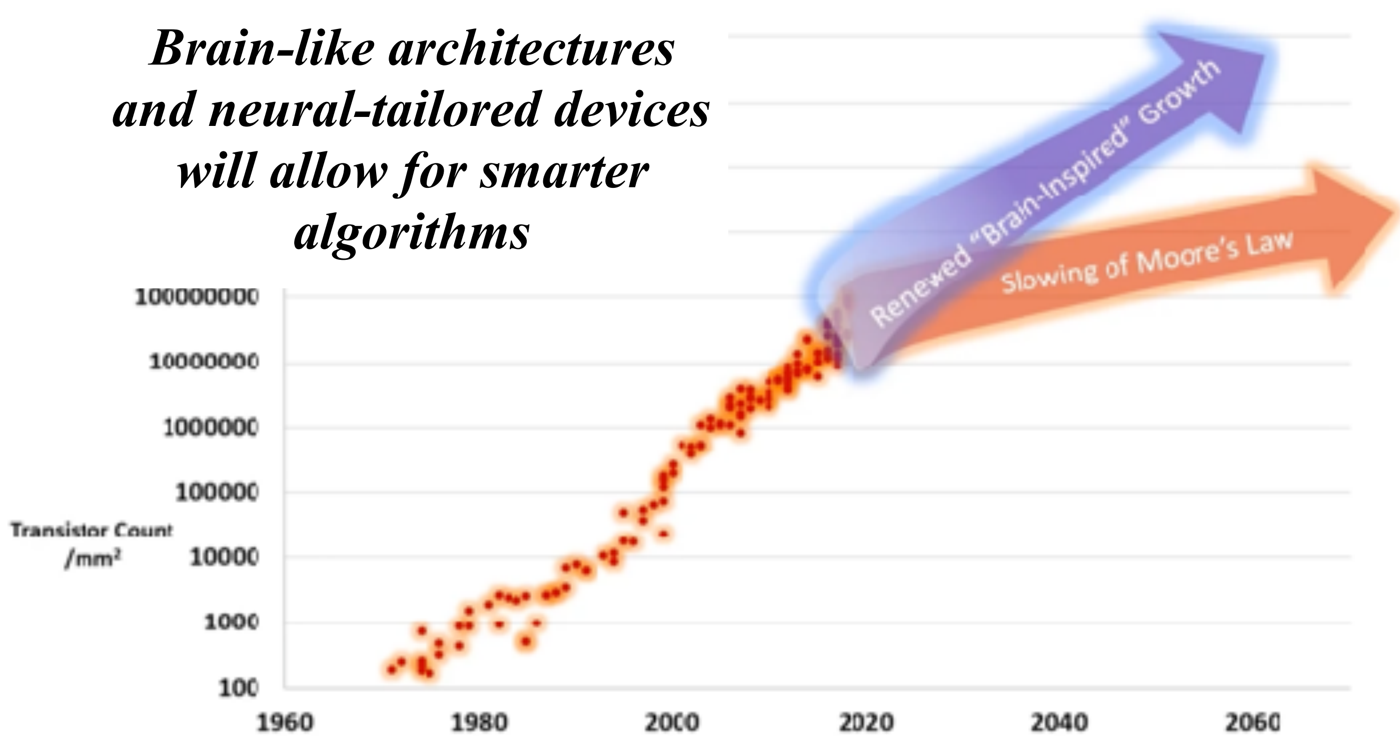

Spiking neural network learns quantum entanglement

***BrainScaleS-2 device**: 512 LIF neurons,*
configurable interaction-weight matrix.
***Communication through spikes**.*

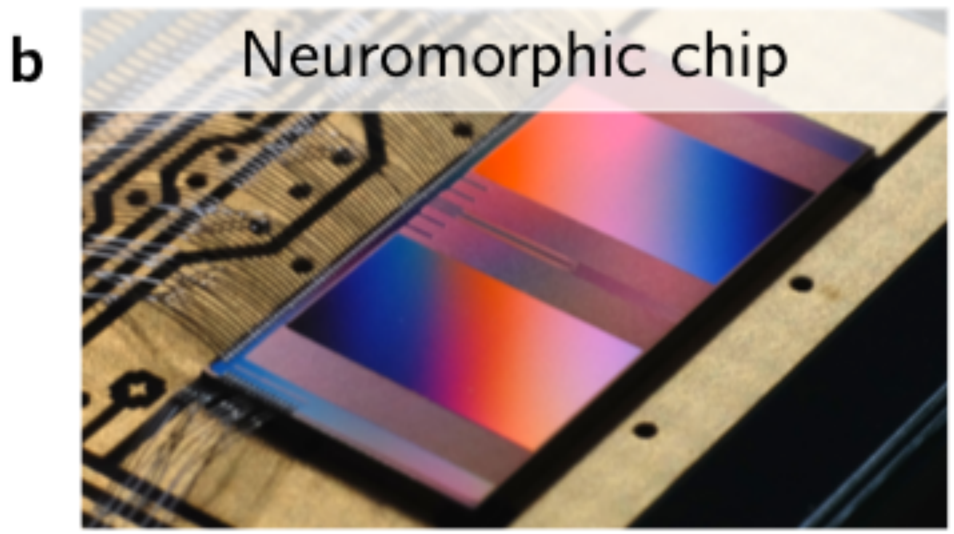

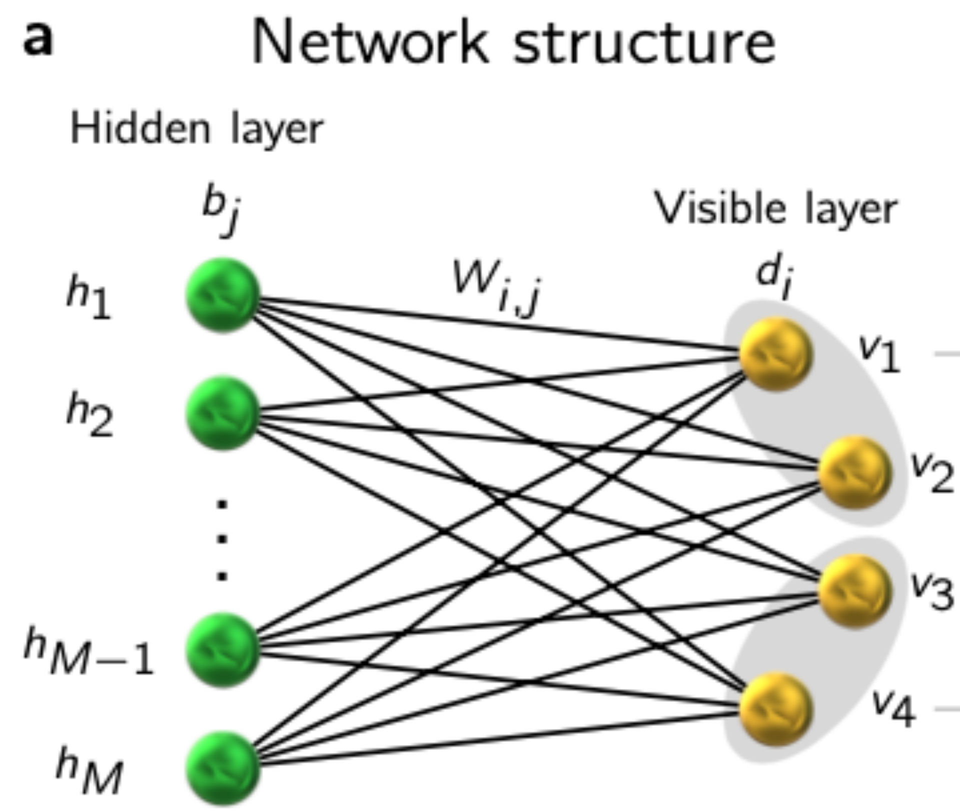

- **Approximate MC Markov chains**, albeit with **different dynamics** from standard methods.
- Accelerated **analog** circuit dynamics: **faster sampling** than with von-Neumann computers.
- **Parallel nature** brings scaling benefits: **better dependence of mixing time with network size**.
- **Neuromorphic RBM network** successfully encodes **GHZ states of *N*=4 qubits**!

Part II´s Conclusions:

- Idea: sequentially learn the updated state's distribution gate by gate
  1. Transformer generative model simulates non-trivial quantum circuits;
  2. Concept demonstration only, efficiency optimisations in place.

J. Carrasquilla, D. Luo, F. Pérez, A. Milstead, B. K. Clark, M. Volkovs, and LA, **Probabilistic Simulation of Quantum Circuits with the Transformer**, arXiv: 1912.11052.

- Exciting possibility: generative modelling with neuromorphic architectures
  1. Proof-of-principle encoding of genuine 4-partite entangled states;
  2. Exploit sampling advantages for complex quantum simulations?

S. Czischek, A. Baumbach, S. Billaudelle, B. Cramer, L. Kades, J. M. Pawlowski, M. Oberthaler, J. Schemmel, M. A. Petrovici, T. Gasenzer, and M. Gärttner, **Spiking neuromorphic chip learns entangled quantum states**, arXiv: 2008.01039.

Part III: application to NISQ computations (work in progress)

The SWAP test: a crucial quantum computing primitive

The “equality algorithm”:

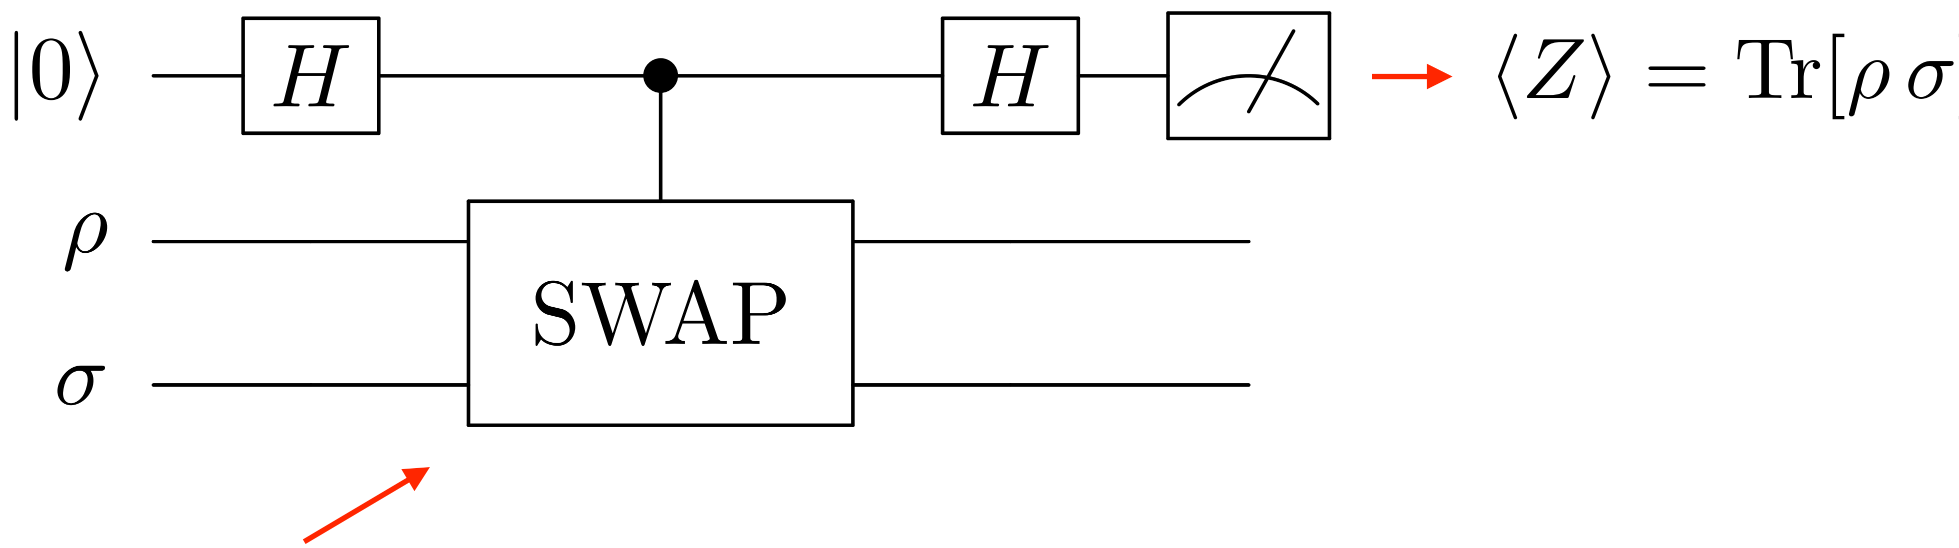

H. Buhrman, R. Cleve, J. Watrous, R. de Wolf, Quantum fingerprinting, Phys. Rev. Lett. **87**, 167902 (2001).

- Directly estimates **overlap between unknown states**.
- Helmstrom **minimum-error state discrimination** without state descriptions (only from samples).
- Elementary procedure **useful for quantum machine learning** algorithms.

An innocent-looking circuit…

Single-qubit SWAP test:

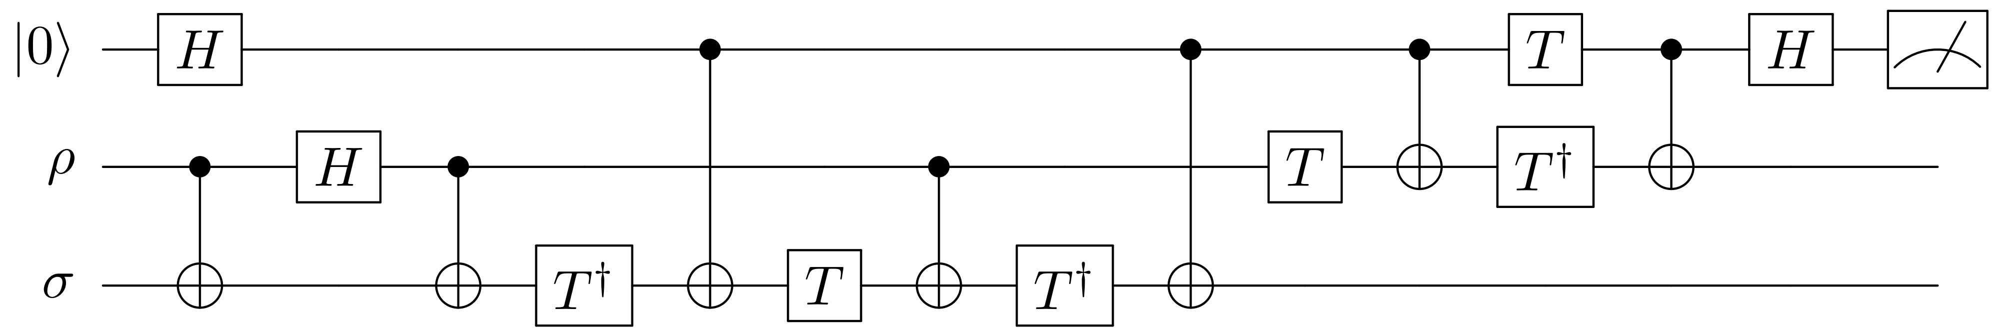

(on IBM´s 5-qubit quantum processor: depth 14 with 7 CNOTs)

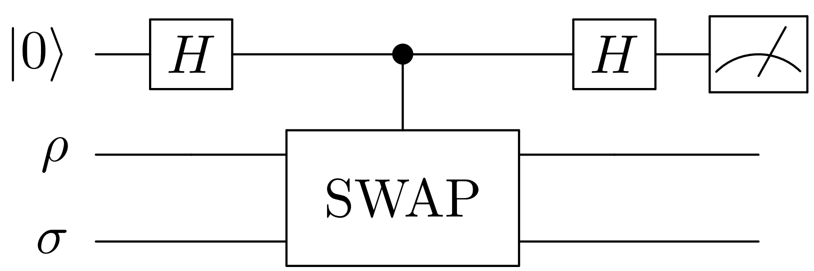

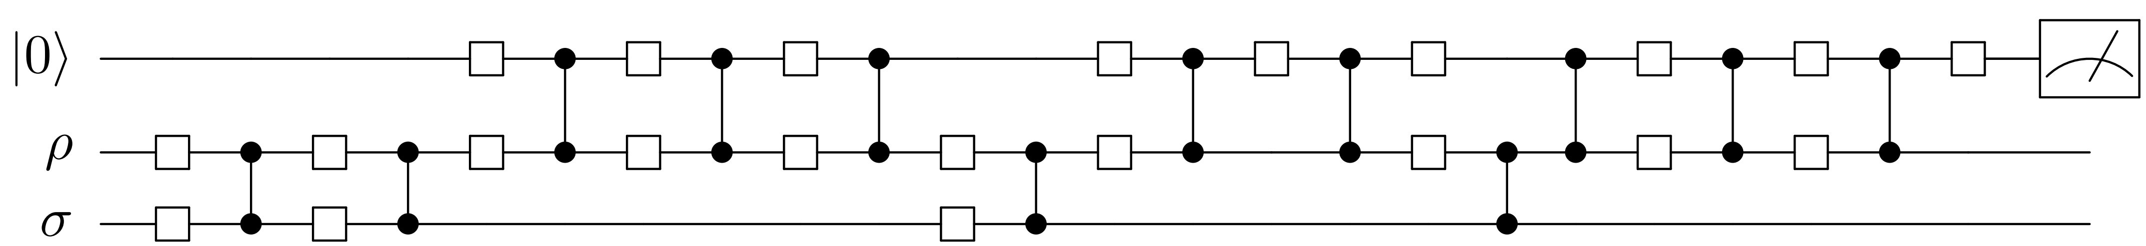

(on Rigetti´s 19-qubit processor: depth 24 with 12 C-PHASEs)

L. Cincio, Y. Subaşı, A. T. Sornborger, and P. J. Coles, New J. Phys. **20**, 113022 (2018).


The NISQ mantra
Run simple things classically and save quantum hardware for the complex stuff

Circuit optimisation for the overlap algorithm

Ancilla-based algorithm:

$U = T^\dagger H$

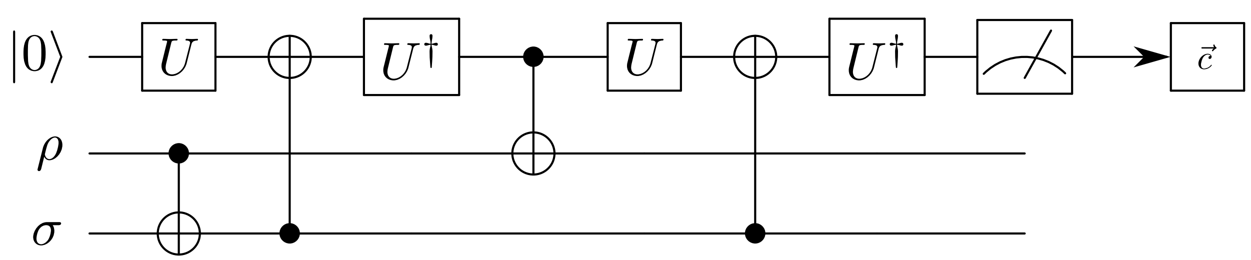

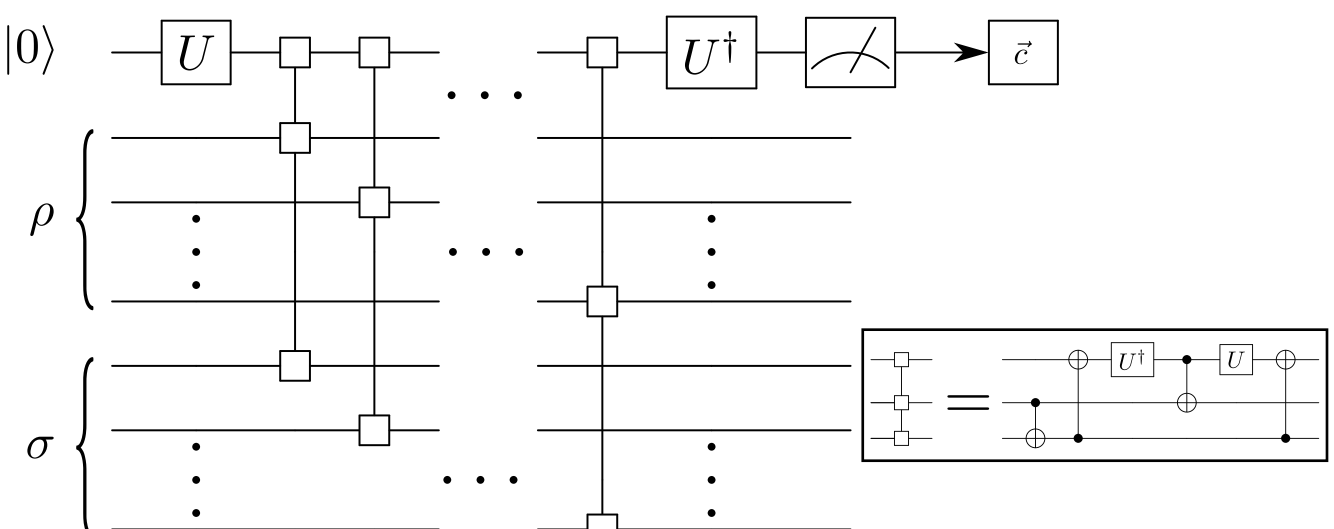

(4N non-nearest-neighbour CNOTs)

L. Cincio, Y. Subaşı, A. T. Sornborger, and P. J. Coles, New J. Phys. **20**, 113022 (2018).

Bell-measurement algorithm:

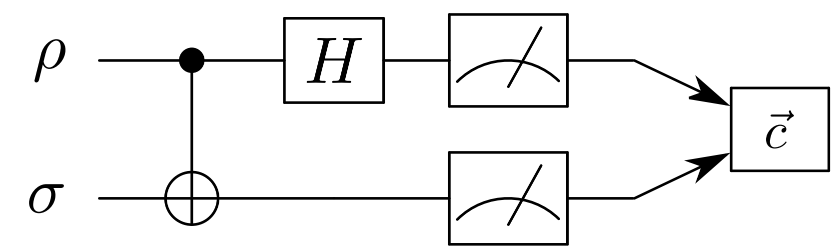

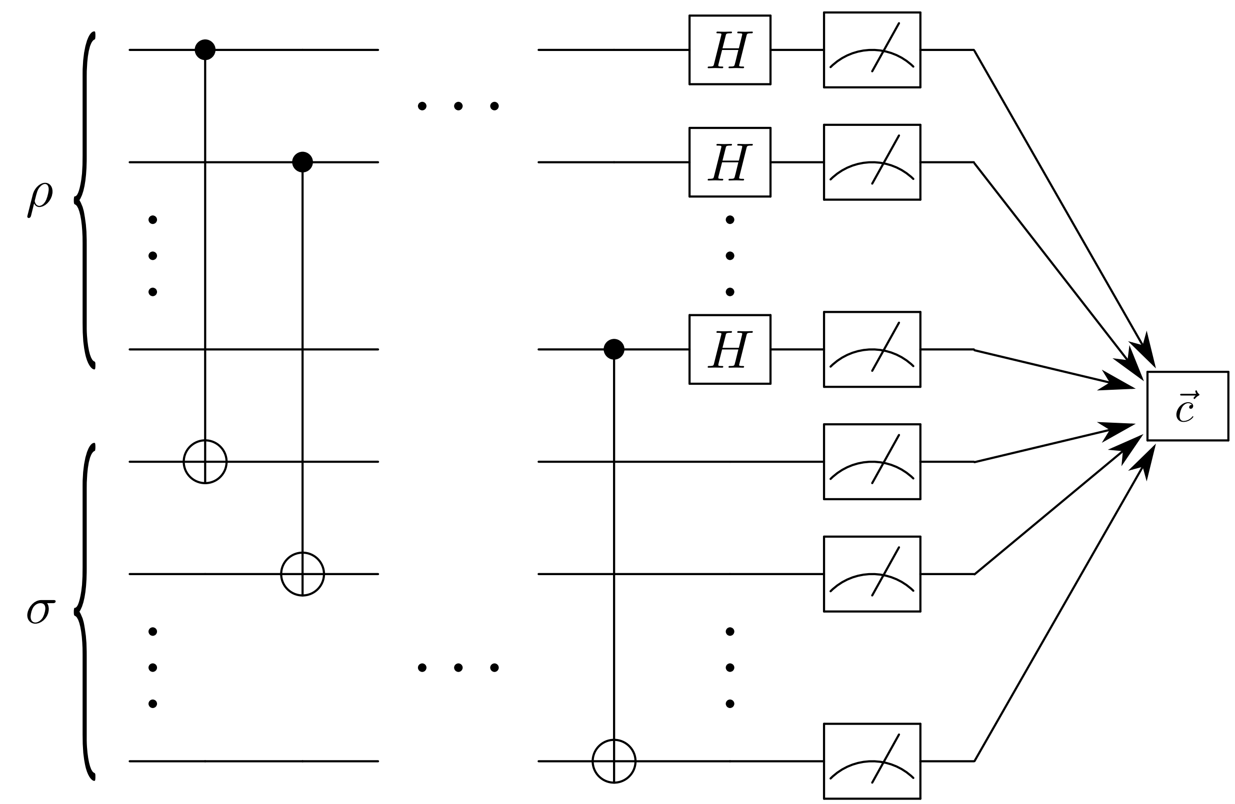

(N non-nearest-neighbour CNOTs)

J. C. García-Escartín and P. Chamorro-Posada, Phys. Rev. A **87**, 052330 (2013).

Idea: estimate overlap from local measurements alone

L. Guerini, J. Carrasquilla, and L. Aolita, **Direct state-overlap estimation without entangling gates**, in preparation.

Statistical overhead due to variance divergence or noisy-gate overhead due to limited connectivity?

Preliminary numerical results

Entanglement-free method versus Bell-measurement one

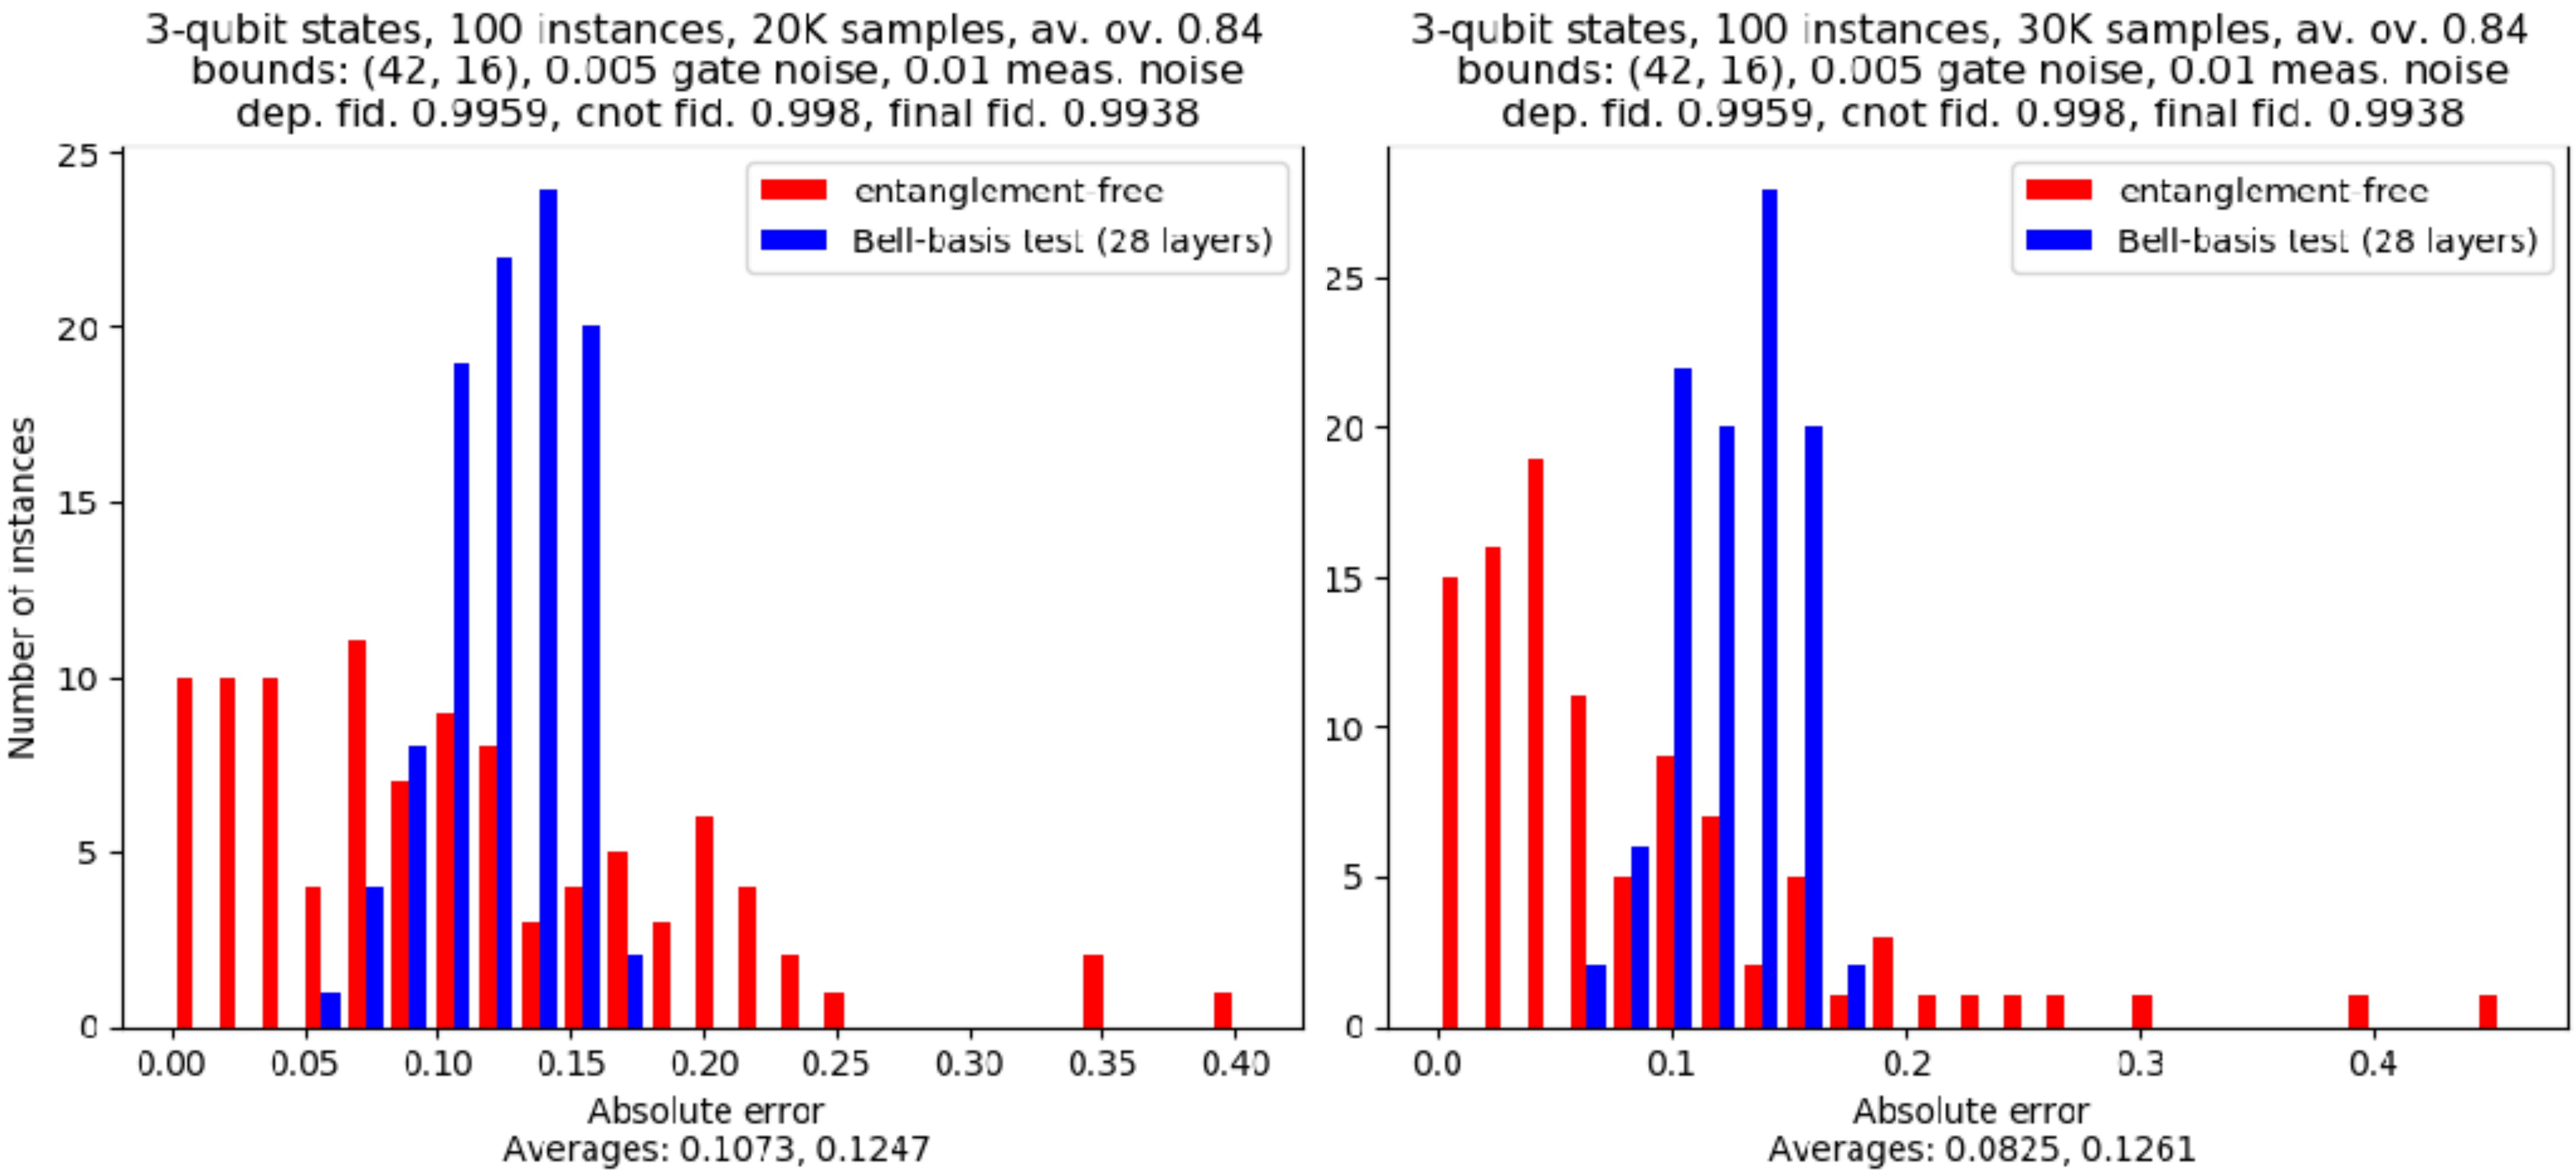

- For fixed number of samples, **statistical error** of entanglement-free method **grows with** ***N*****.**
- But **systematic error of** coherent tests with **noisy entangling gates also grows with** *N*.

Part III´s Conclusions:

- Direct state-overlap estimation without entangling gates:

  1. Both systems need not be simultaneously available.

  2. The noisier the gates, the more convenient the method.

L. Guerini, J. Carrasquilla, and L. Aolita, **Direct state-overlap estimation without entangling gates**, in preparation.

Global conclusions:

- Probabilistic formulation of quantum statics and dynamics amenable for generative machine learning:
    1. Record breaking state and process tomography
    2. Simulations of quantum circuits with local gates.
    3. Swap test without entangling gates.

Forwards

Hybrid classical-quantum algorithms

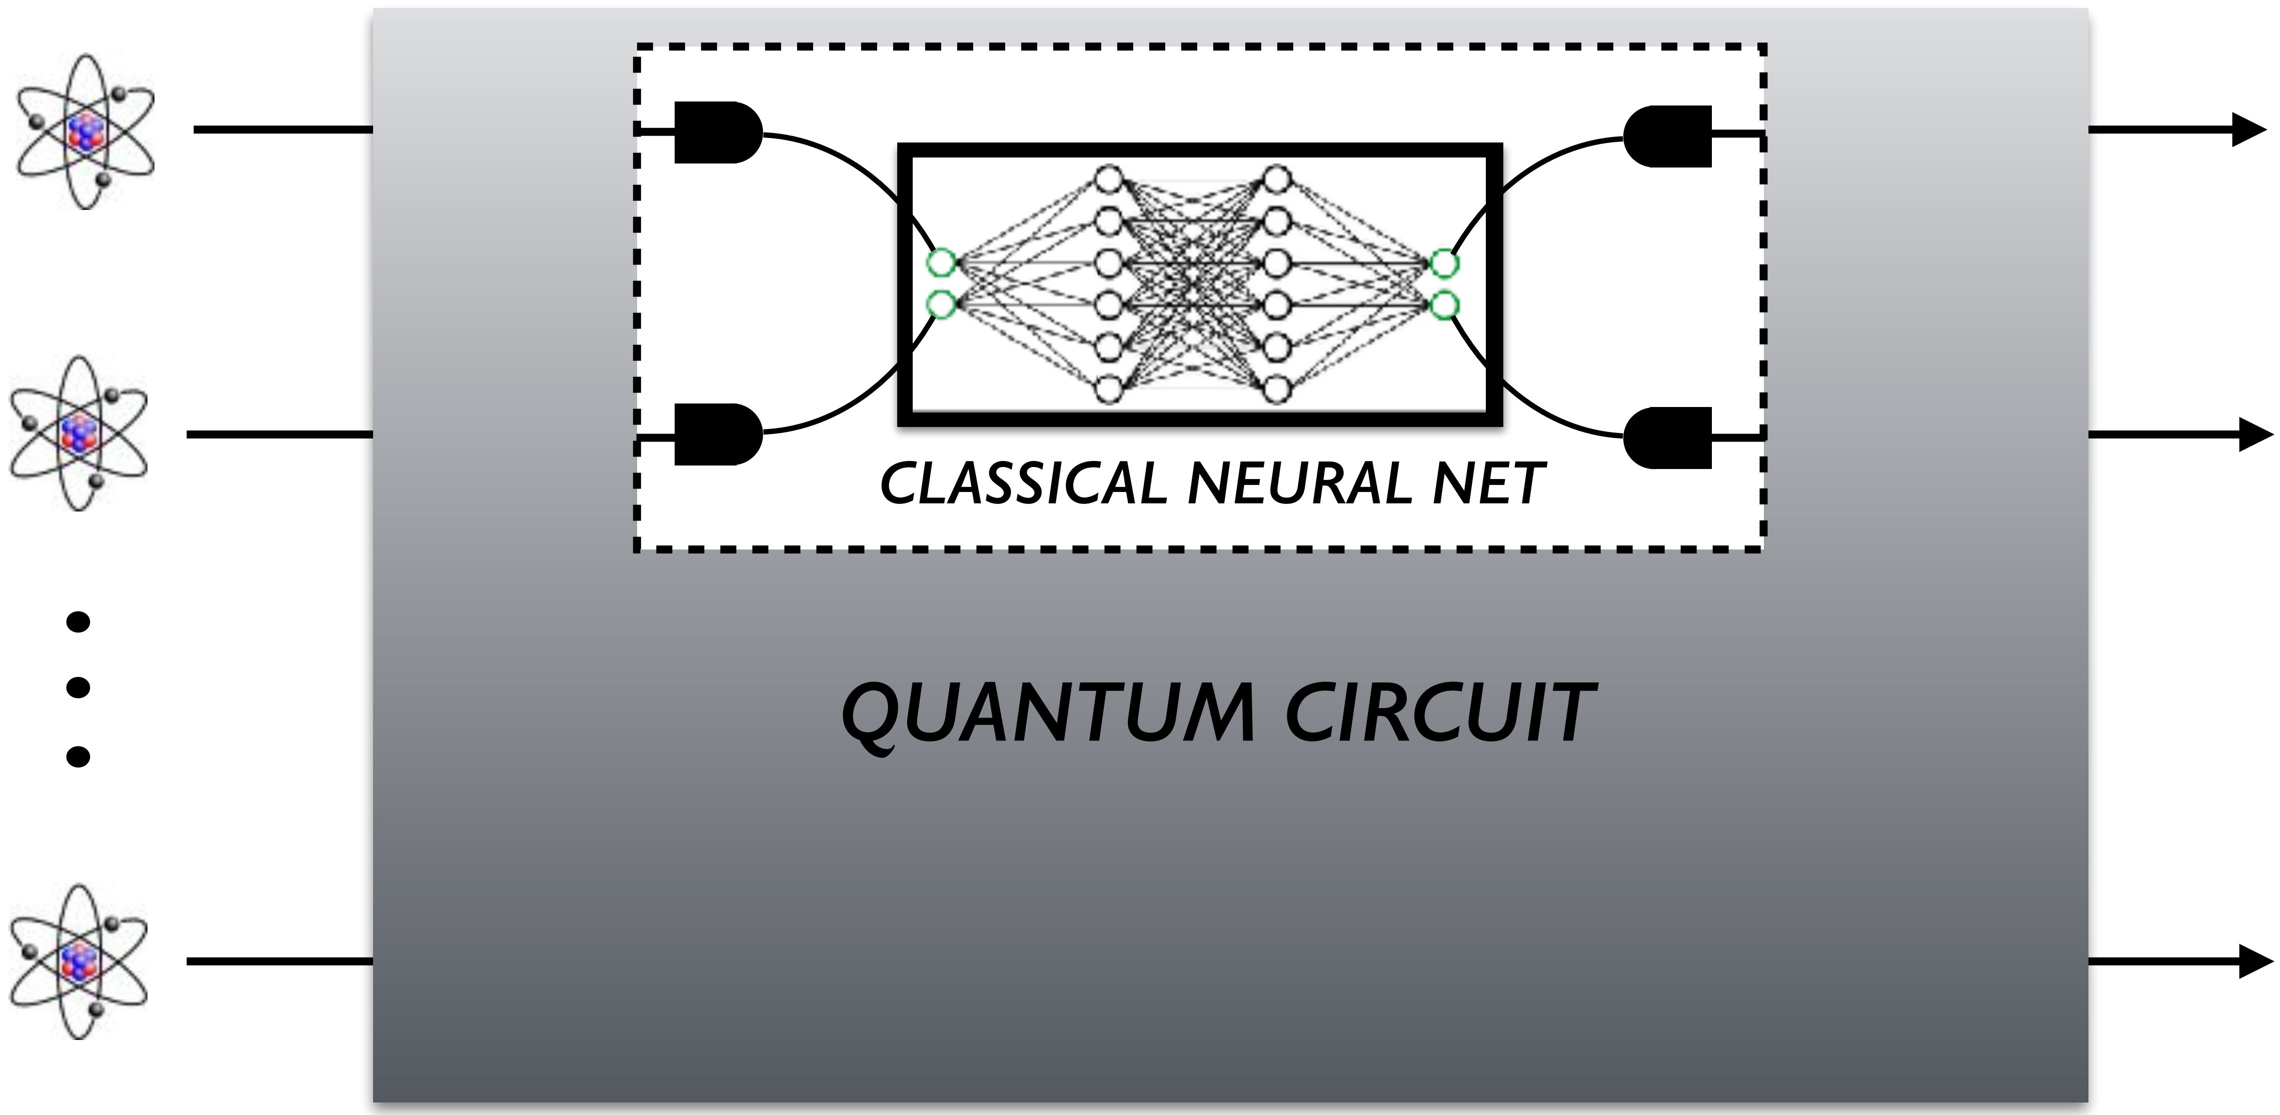

Potential applications:

- Quantum artificial intelligence
- Complex combinatorial optimisations
- Quantum chemistry
- Large differential-equation systems

F. de Melo e L. Aolita, Ciencia Hoje **360** (2019)

Main external collaborators:

Juan Carrasquilla
Vector Institute (Toronto)

Giacomo Torlai
Flatiron Institute (NYC)
AWS (Pasadena)

Roger Melko
Perimeter Institute (Waterloo)

Chris. J. Woods
IBM Quantum (NYC)

Giuseppe Carleo
Flatiron Institute (NYC)
EPFL (Lausanne)

Family at Rio:

Leonardo Guerini
Postdoc

Thais de Lima Silva
Postdoc

Eric G. A. Cavalcanti
PhD student

Renato M. S. Farias
PhD student

Lucas A. Borges
PhD student

Former:

Marcio M. Taddei
Postdoc

Carolina Gigliotti
Masters student

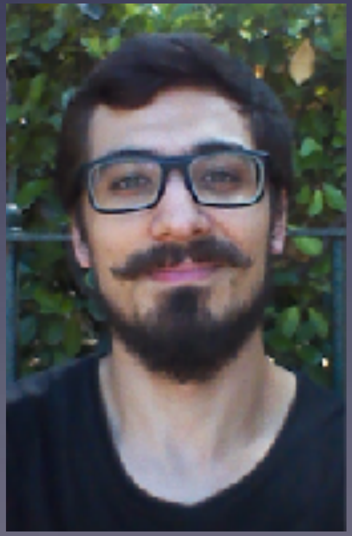

Ranieri V. Neri
Postdoc

But, what about Abu Dhabi?

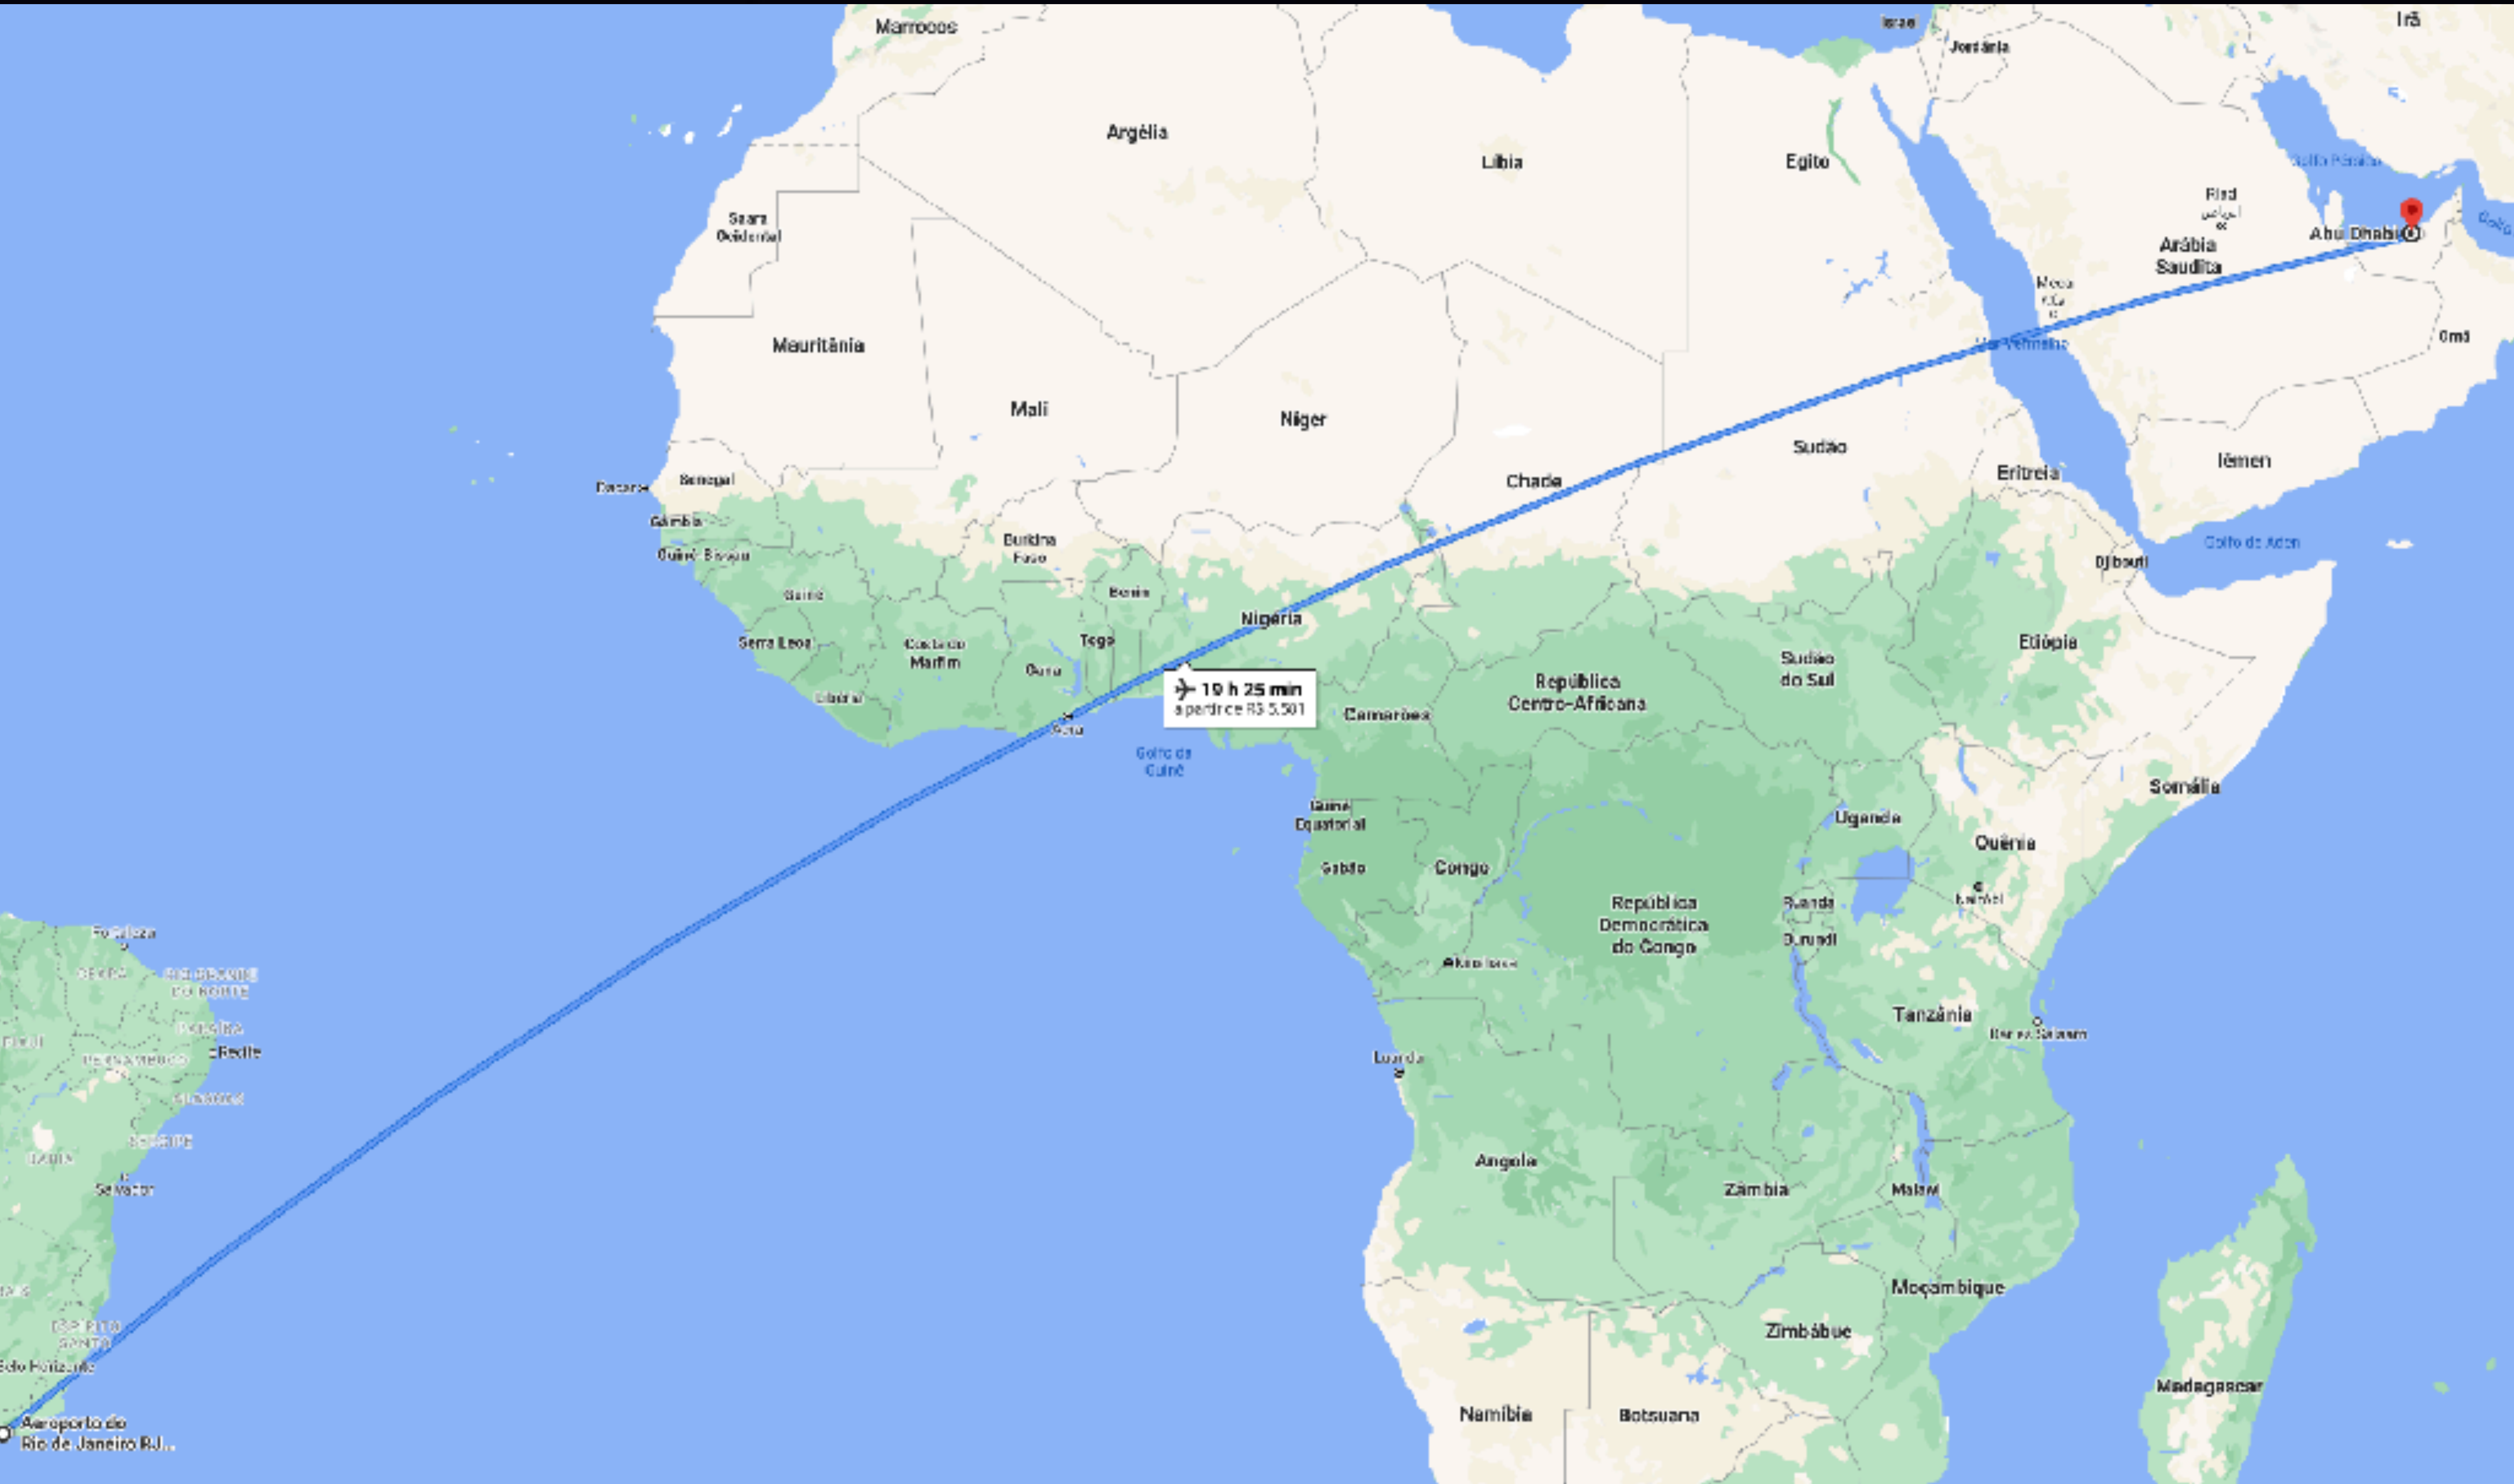

Technology
Innovation
Institute

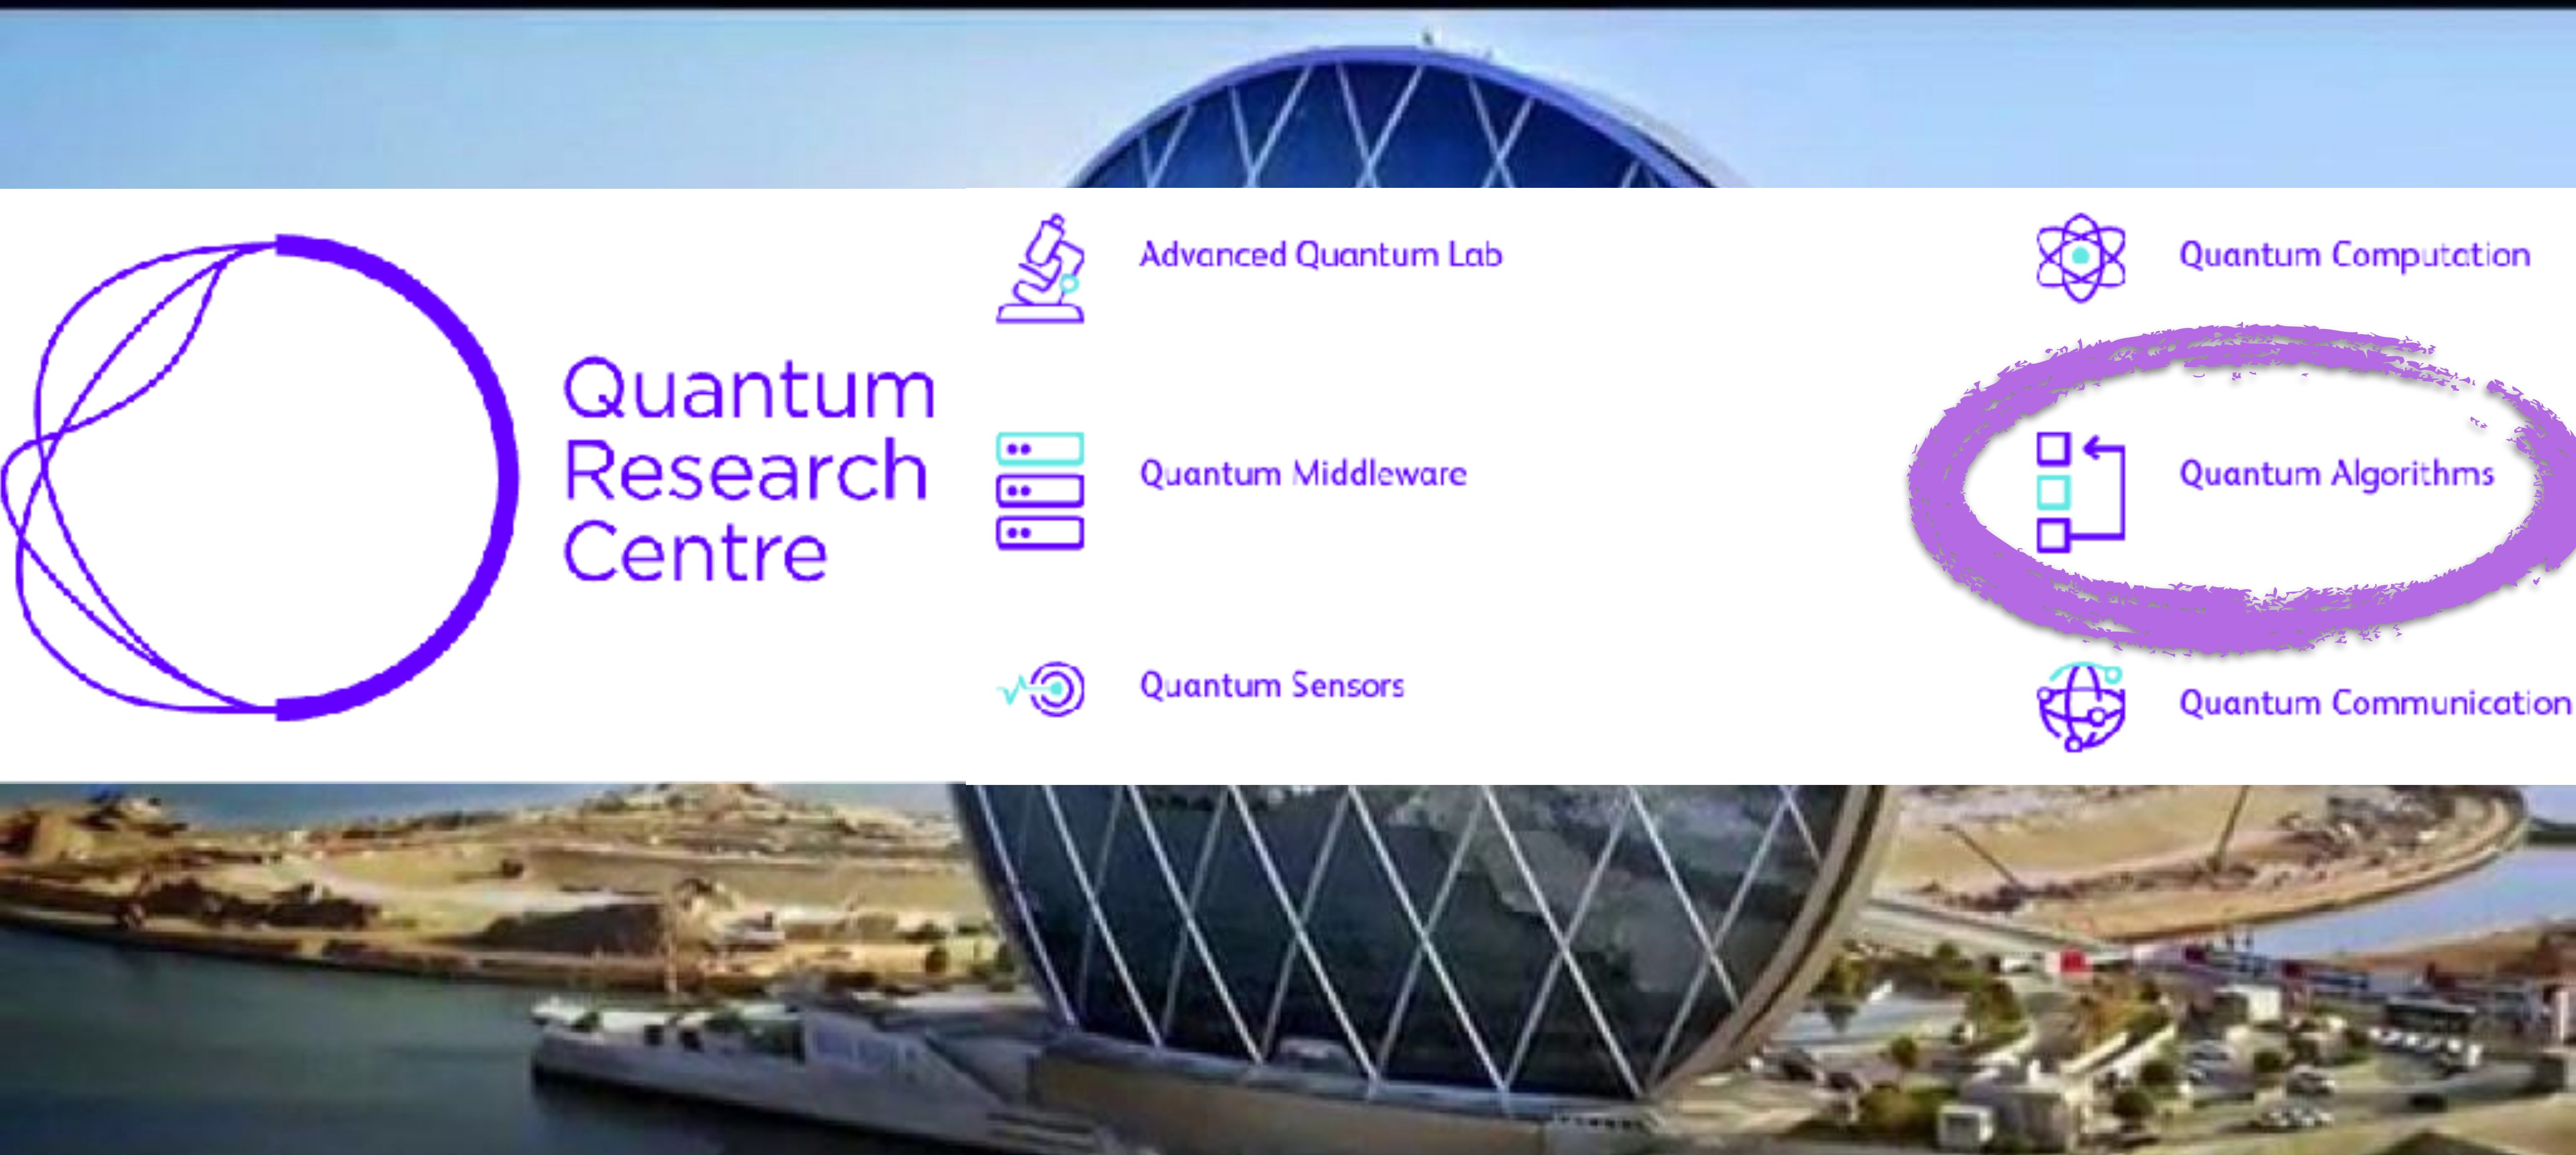

An excellence sci + tech hub in the middle east

What about Abu Dhabi?

A super exciting new endeavour!

(Positions open!)

J. Carrasquilla, G. Torlai, R. Melko, and L. Aolita, **Reconstructing quantum states with generative models**, Nature Machine Intelligence **1**, 155 (2019).

G. Torlai, C. J. Wood, A. Acharya, G. Carleo, J. Carrasquilla, and LA, **Quantum process tomography with unsupervised learning and tensor networks**, arXiv: 2006.02424.

J. Carrasquilla, D. Luo, F. Pérez, A. Milstead, B. K. Clark, M. Volkovs, and L. Aolita, **Probabilistic Simulation of Quantum Circuits with the Transformer**, arXiv: 1912.11052 (2020).

L. Guerini, J. Carrasquilla, and L. Aolita, **Direct state-overlap estimation without entangling gates**, in preparation.

F. de Melo e L. Aolita, Ciencia Hoje **360** (2019).

Thanks