Quantum Machine Learning

Adam Brown, HEP-AI

Quantum Computing

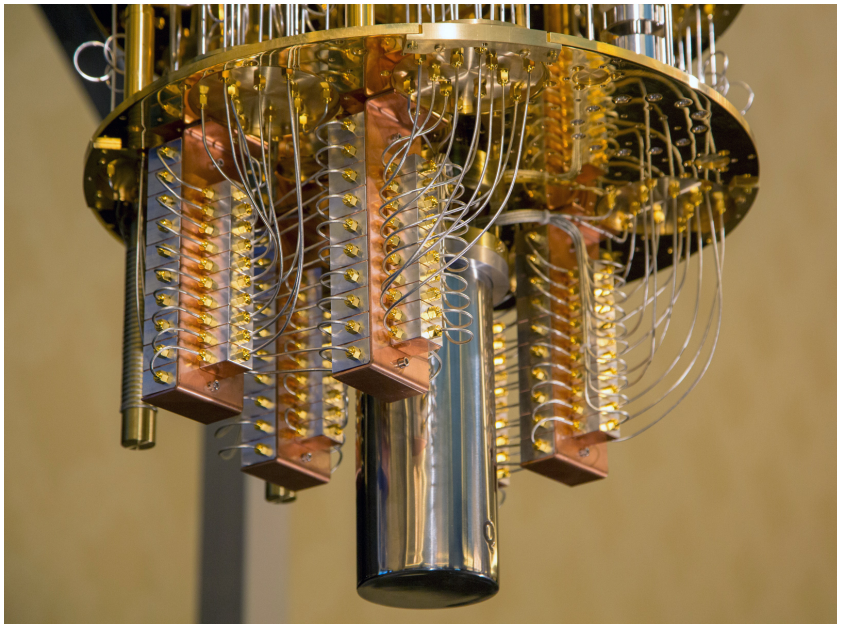

+

Machine Learning

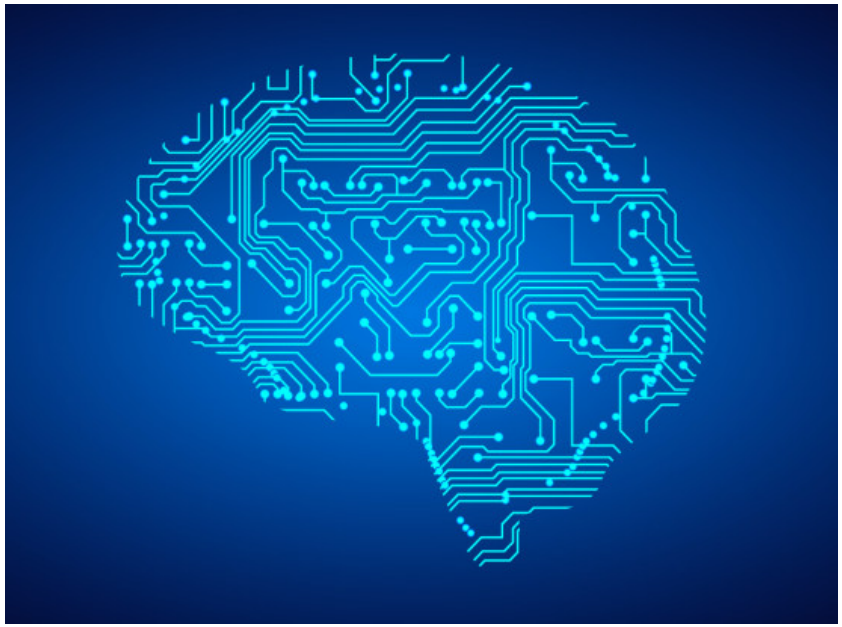

Quantum Machine Learning

Quantum Computing

so hot

+

Machine Learning

so so hot

Quantum Machine Learning

Quantum Machine Learning

“I would predict that in 10 years there’s nothing but quantum machine learning.”
[famous QML researcher in 2015]

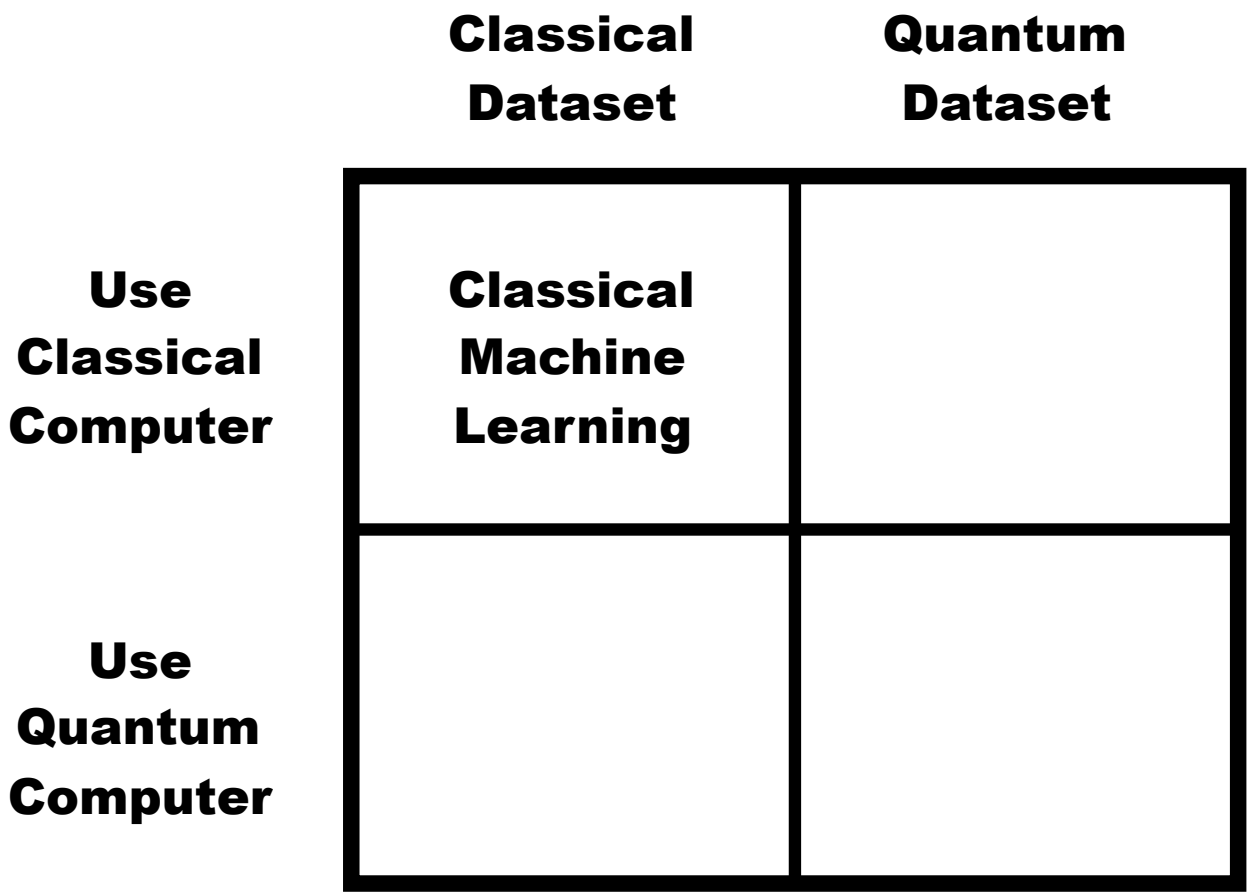
Classical Dataset
Quantum Dataset
Use Classical Computer
Use Quantum Computer
Classical Machine Learning

The Power of Quantum Computers (in theory)

The Power of Quantum Computers (in theory)

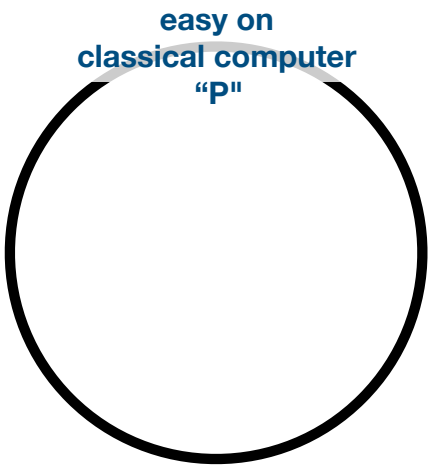

The Power of Quantum Computers (in theory)

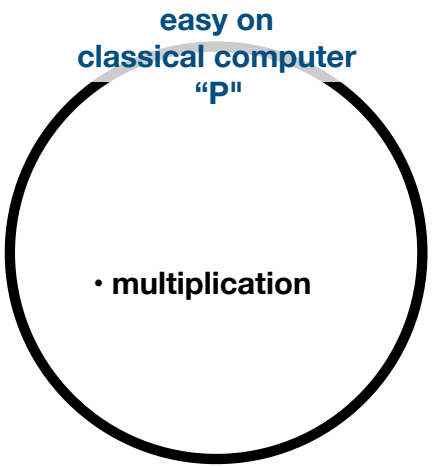

The Power of Quantum Computers (in theory)

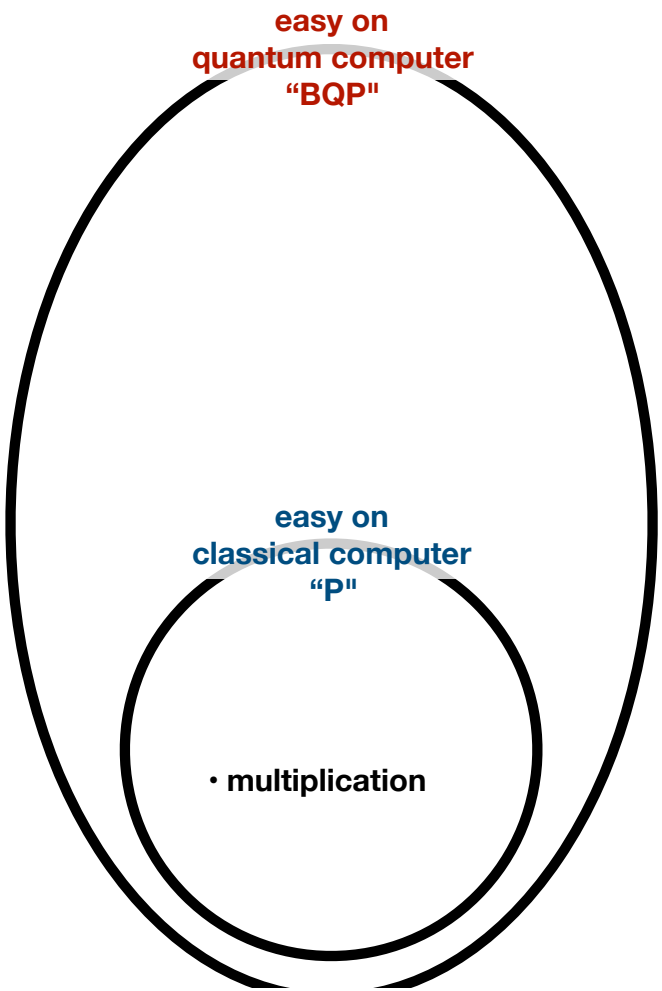

The Power of Quantum Computers (in theory)

(i) BQP contains P

a quantum computer can (quickly) do
everything a classical computer can (quickly) do

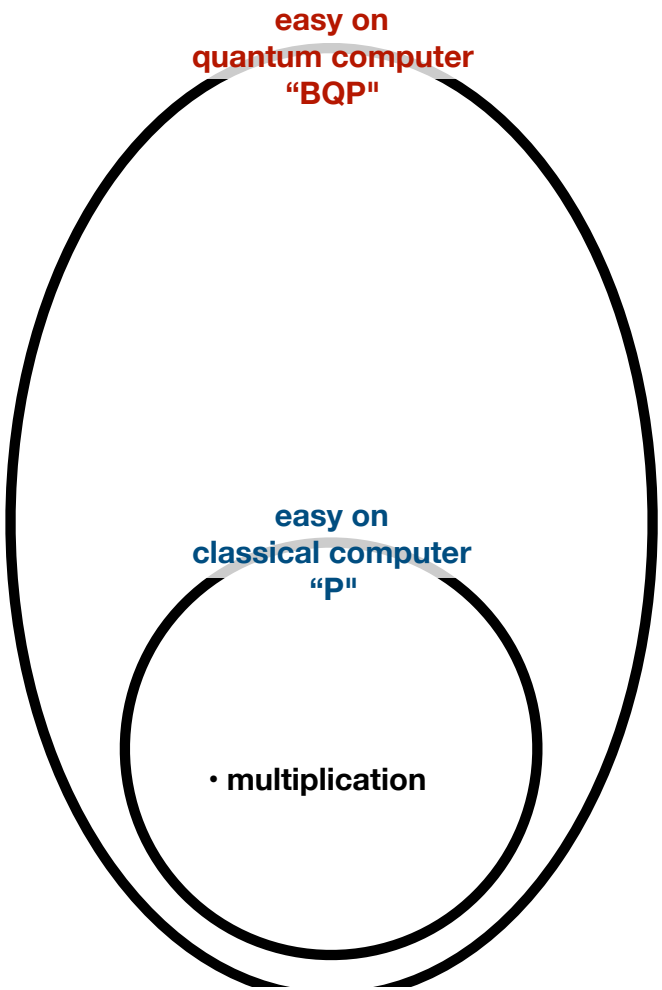

The Power of Quantum Computers (in theory)

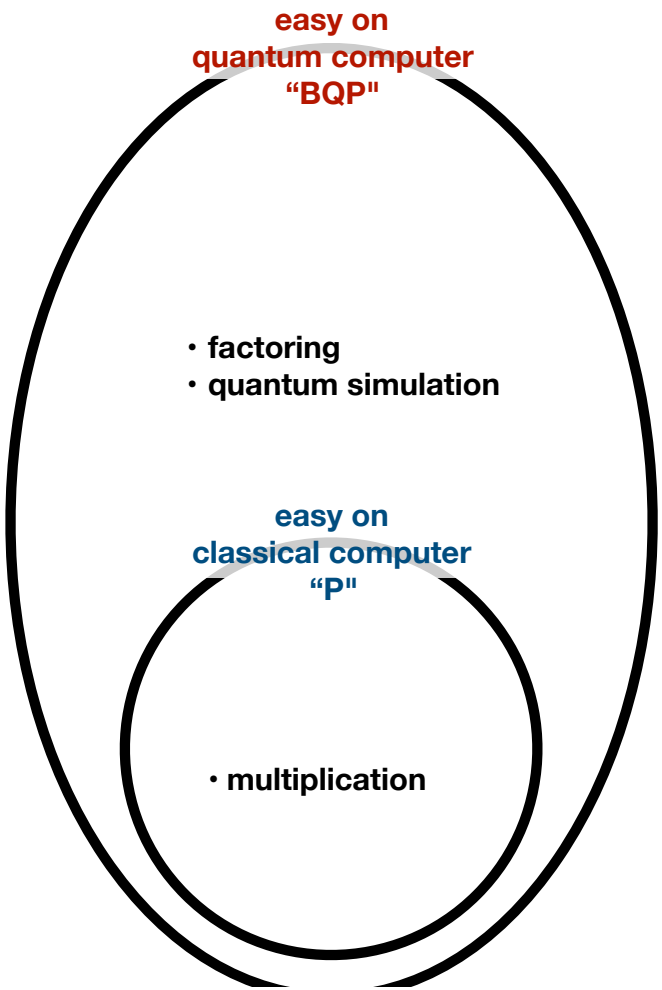

(i) BQP contains P

a quantum computer can (quickly) do
everything a classical computer can (quickly) do

(ii) P does not contain BQP

a quantum computer can do some things quickly
that a classical computer cannot do quickly

↑

ENCOURAGING

The Power of Quantum Computers (in theory)

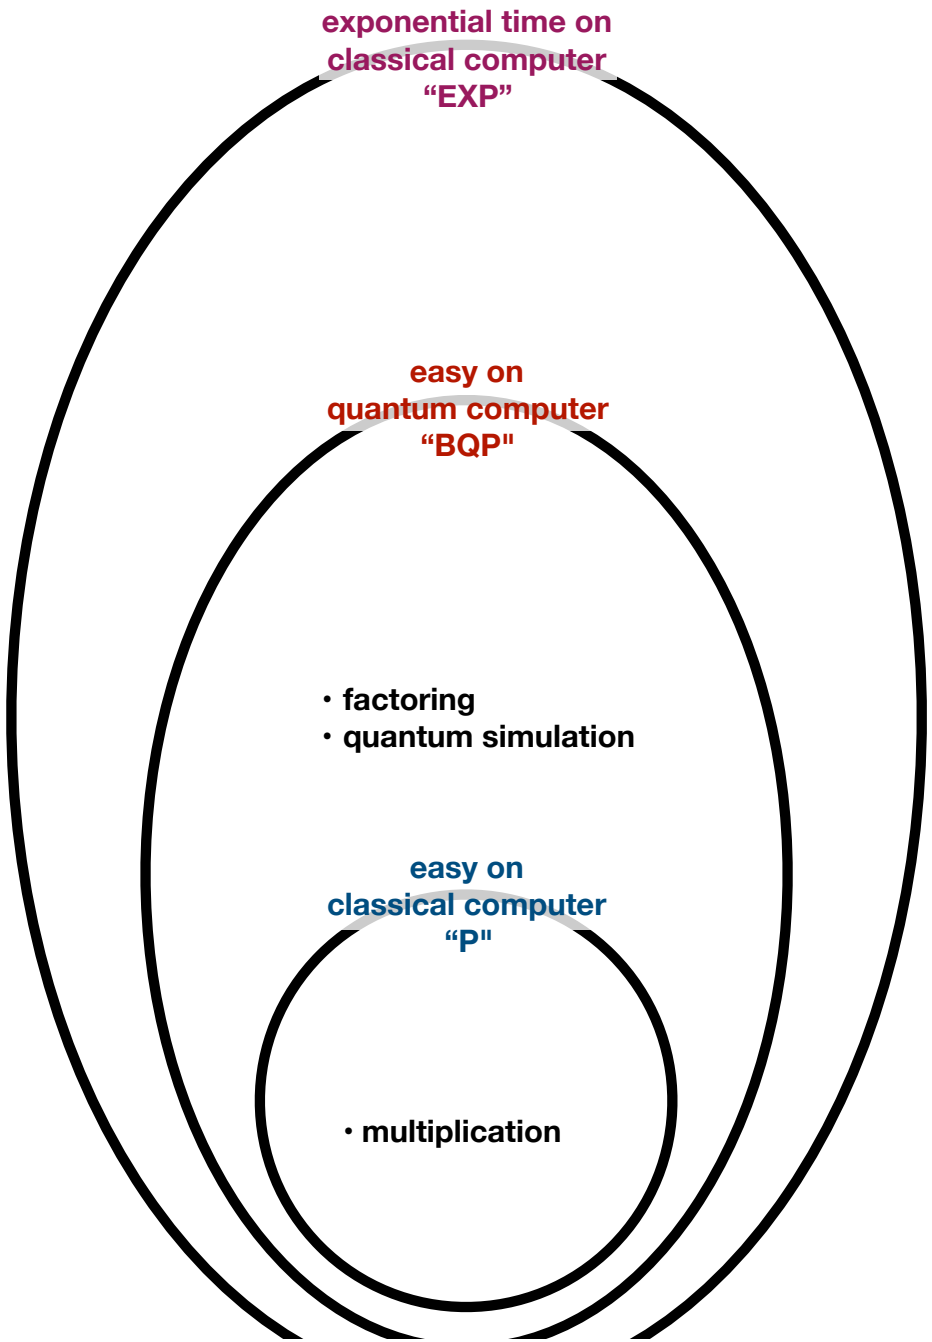

(i) BQP contains P

a quantum computer can (quickly) do
everything a classical computer can (quickly) do

(ii) P does not contain BQP

a quantum computer can do some things quickly
that a classical computer cannot do quickly

↑

ENCOURAGING

(iii) EXP contains BQP

a classical computer can **eventually** do
everything a quantum computer can do

The Power of Quantum Computers (in theory)

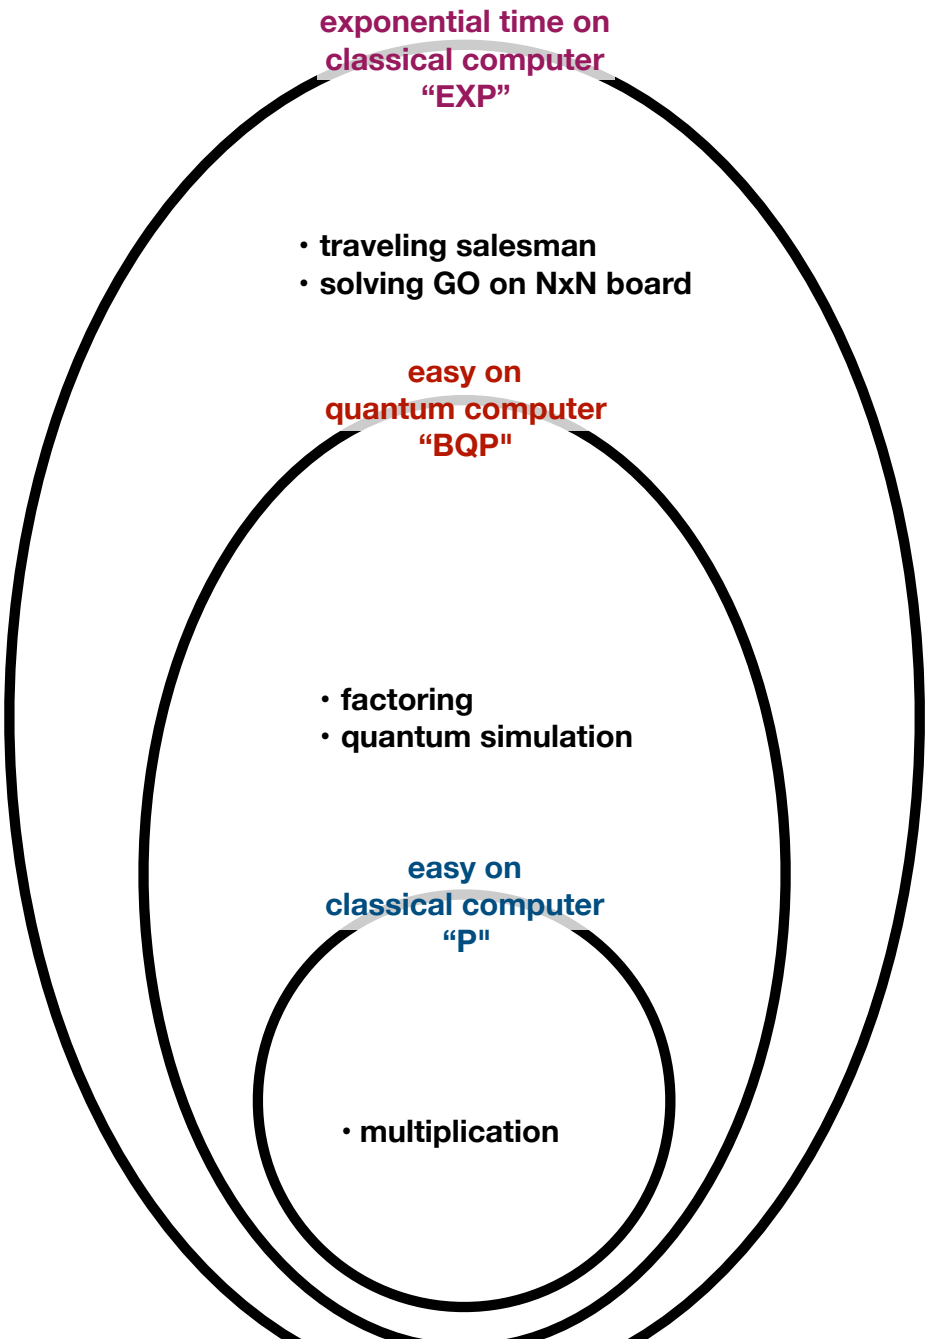

(i) BQP contains P

a quantum computer can (quickly) do
everything a classical computer can (quickly) do

(ii) P does not contain BQP

a quantum computer can do some things quickly
that a classical computer cannot do quickly

↑

ENCOURAGING

(iii) EXP contains BQP

a classical computer can **eventually** do
everything a quantum computer can do

(iv) BQP does not contain EXP

a quantum computer does **not** give
generic exponential speed-up

The Power of Quantum Computers (in theory)

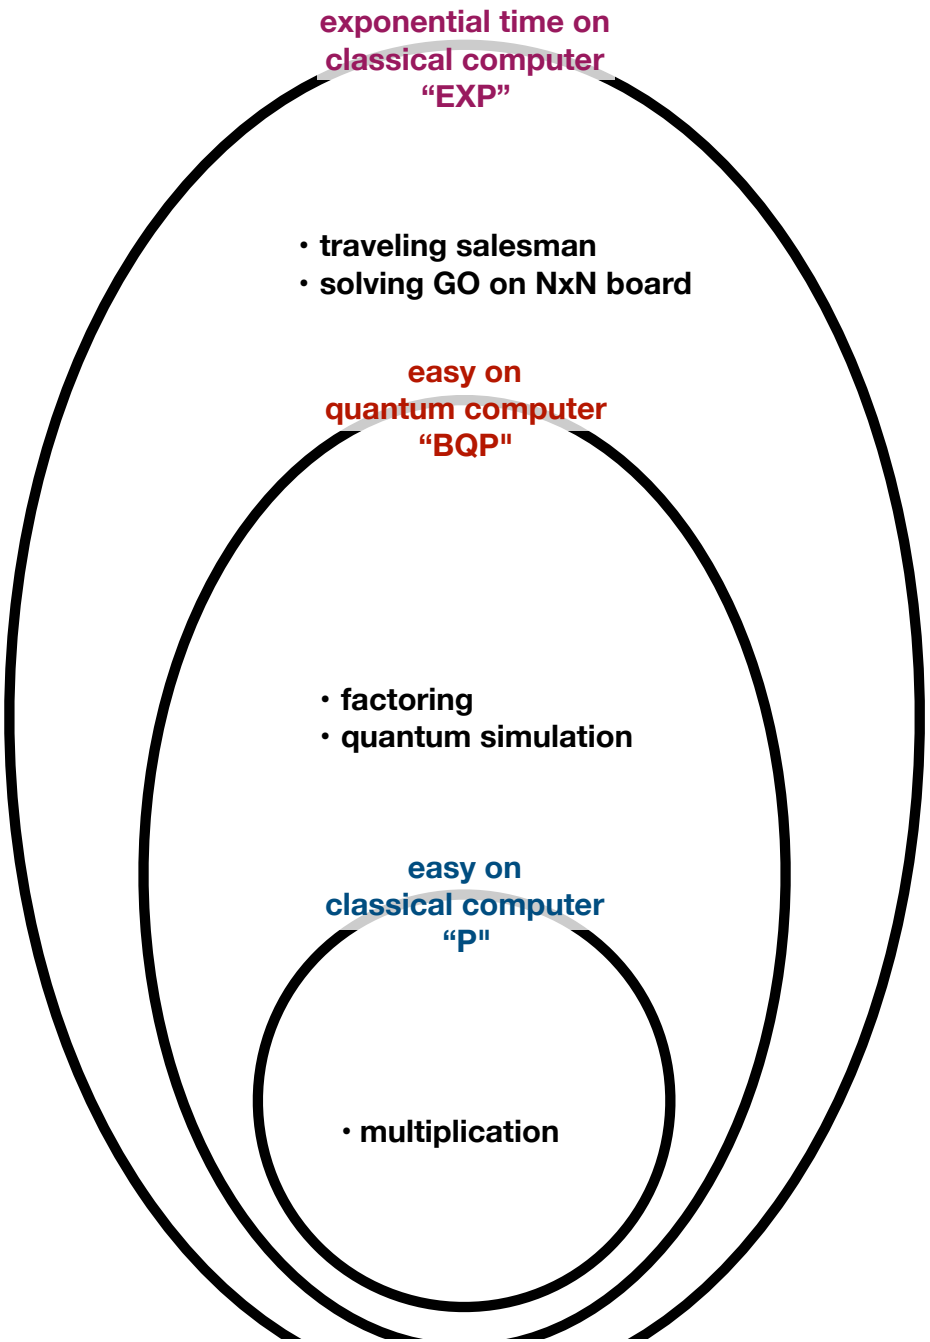

(i) BQP contains P

a quantum computer can (quickly) do
everything a classical computer can (quickly) do

(ii) P does not contain BQP

a quantum computer can do some things quickly
that a classical computer cannot do quickly

↑

ENCOURAGING

(iii) EXP contains BQP

a classical computer can **eventually** do
everything a quantum computer can do

(iv) BQP does not contain EXP

a quantum computer does **not** give
generic exponential speed-up

↑

DISCOURAGING

The Power of Quantum Computers (in theory)

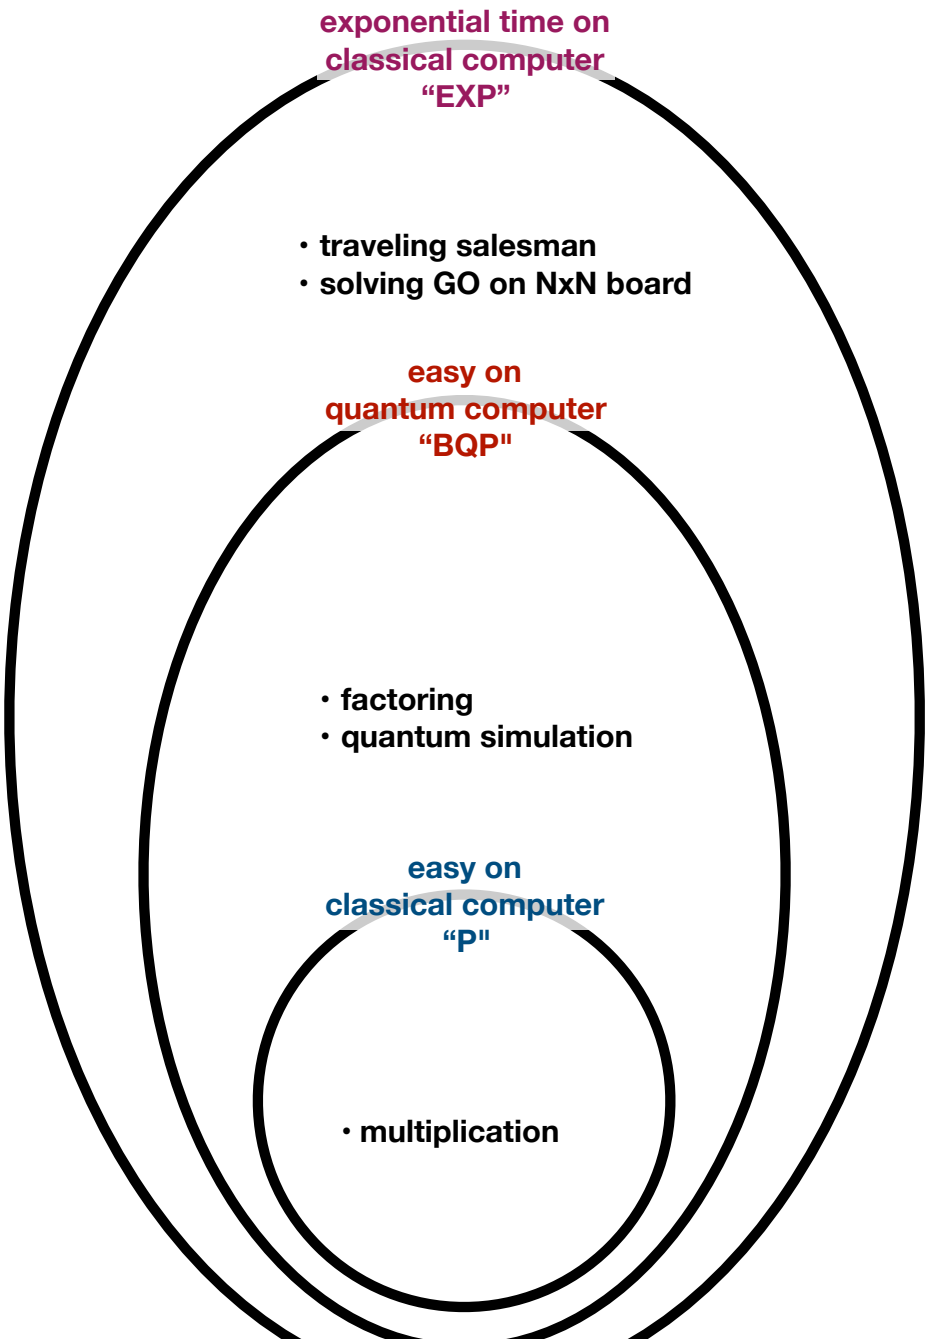

(i) BQP contains P

a quantum computer can (quickly) do
everything a classical computer can (quickly) do

(ii) P does not contain BQP

a quantum computer can do some things quickly
that a classical computer cannot do quickly

↑

ENCOURAGING

(iii) EXP contains BQP

a classical computer can **eventually** do
everything a quantum computer can do

(iv) BQP does not contain EXP

a quantum computer does **not** give
generic exponential speed-up

↑

“NO FREE LUNCH”

The Power of Quantum Computers (in theory)

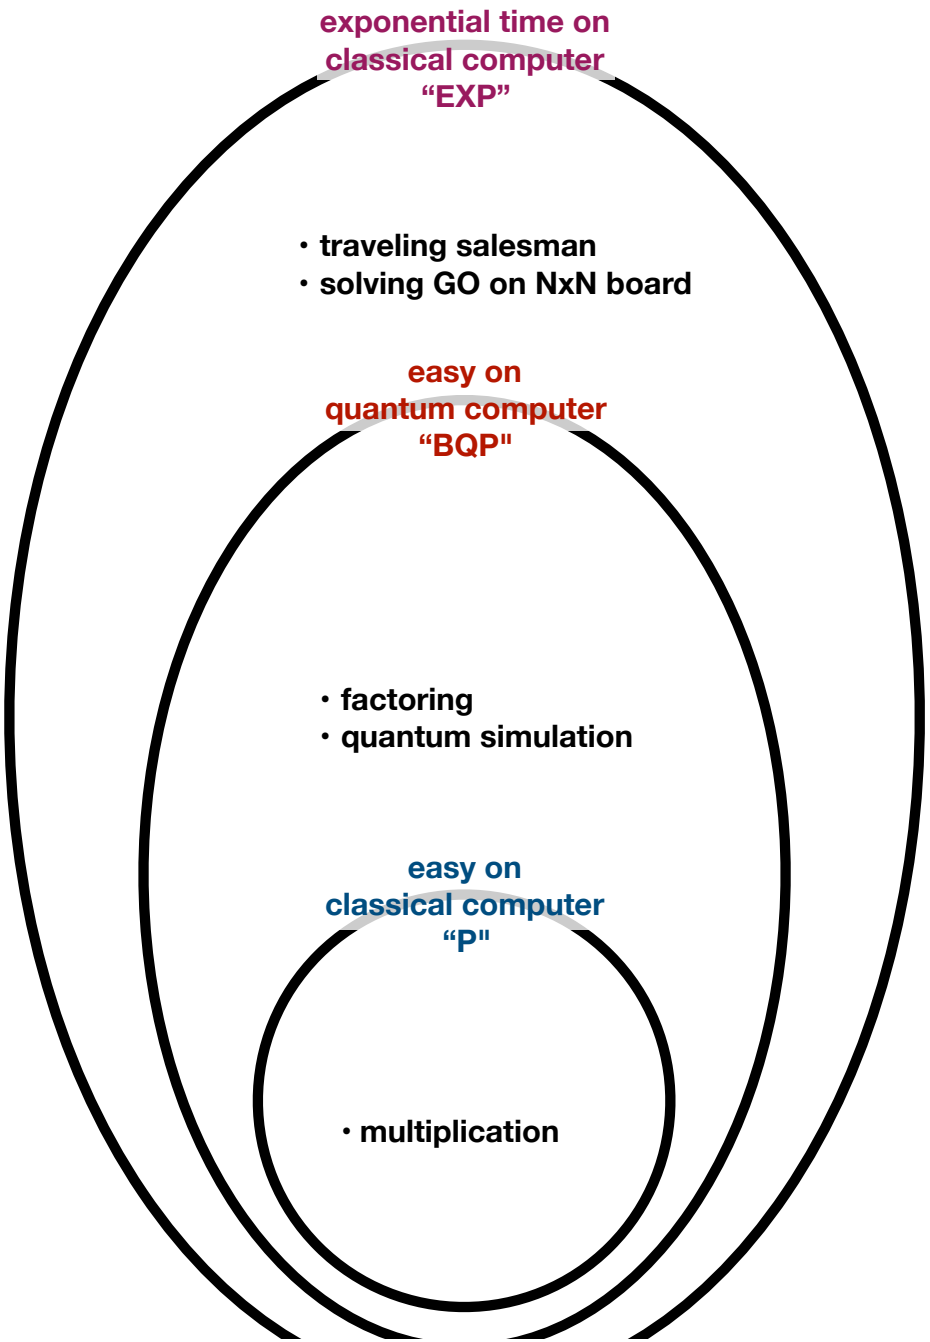

quantum superpower:
EXPONENTIATION

“HILBERT SPACE IS HUGE”

to describe state of *N* qubits takes 2^N real numbers

$$\alpha_1|00000000\rangle + \ldots + \alpha_{2^N}|11111111\rangle$$

↑

ENCOURAGING

(iii) EXP contains BQP

a classical computer can **eventually** do everything a quantum computer can do

(iv) BQP does not contain EXP

a quantum computer does **not** give **generic** exponential speed-up

↑

“NO FREE LUNCH”

The Power of Quantum Computers (in theory)

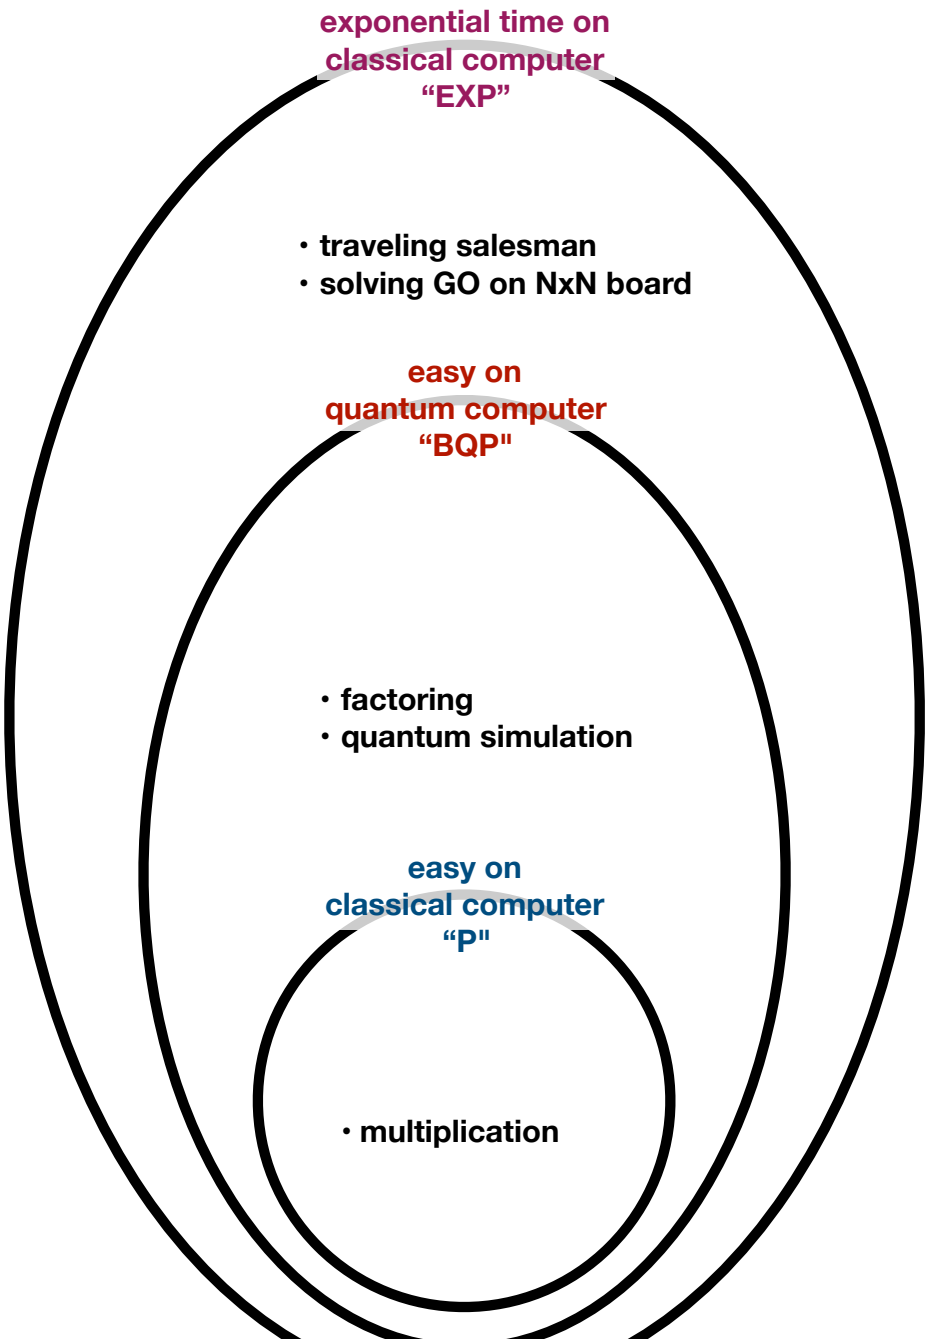

quantum superpower:
EXPONENTIATION

“HILBERT SPACE IS HUGE”

to describe state of *N* qubits takes 2^N real numbers

$$\alpha_1|00000000\rangle + \ldots + \alpha_{2^N}|11111111\rangle$$

↑

ENCOURAGING

quantum superweakness:
LINEARITY

cannot directly see wave function

↑

“NO FREE LUNCH”

Encouraging Fact #1

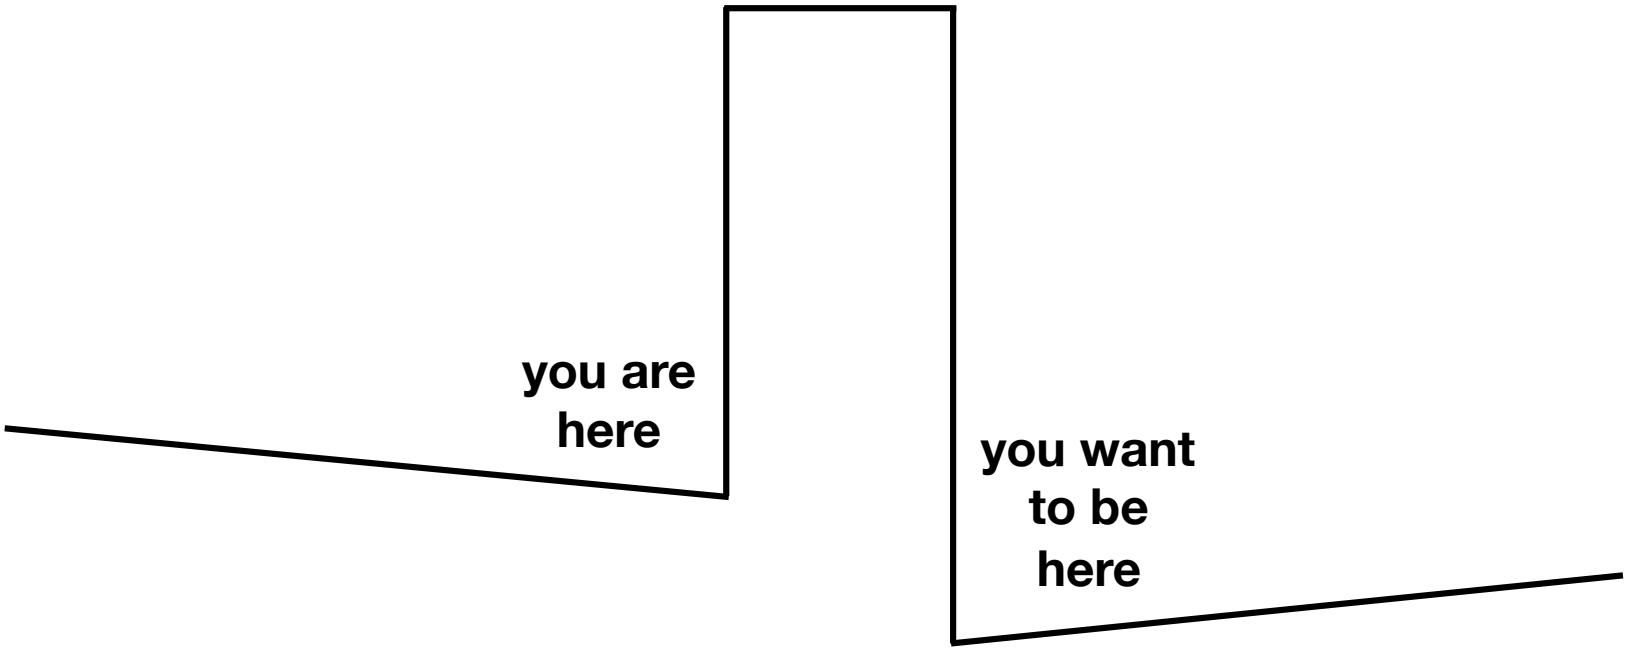

Encouraging Fact #1

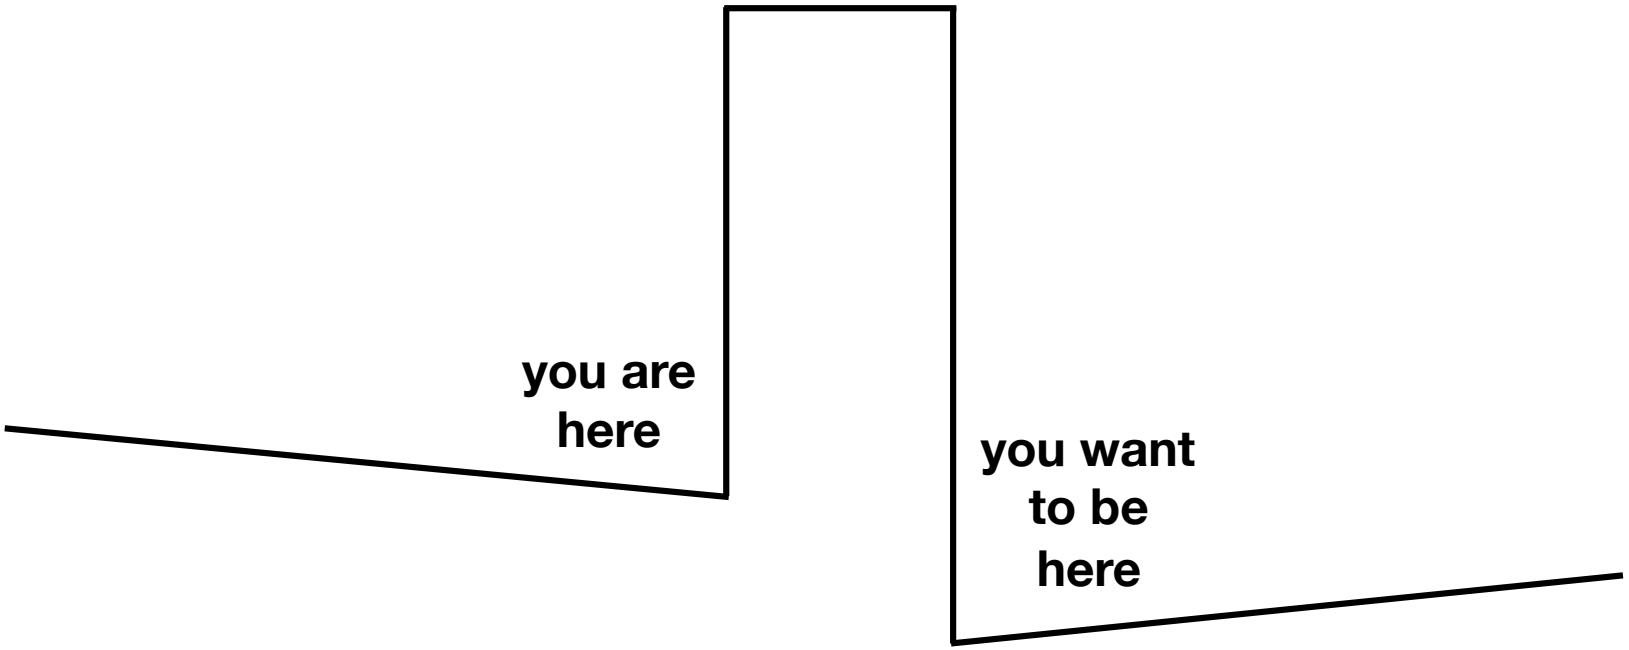

thermally fluctuate $\sim e^{-\text{height}/T}$

Encouraging Fact #1

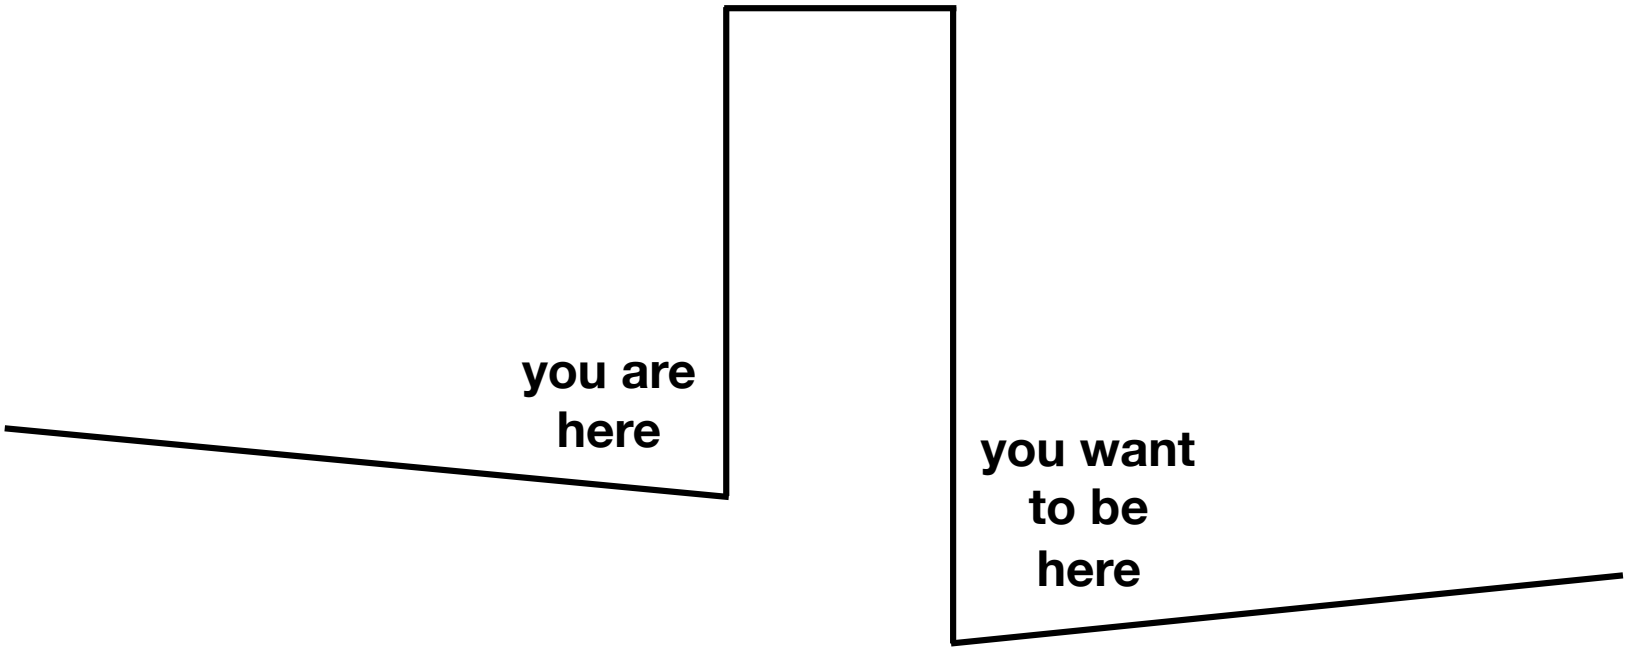

thermally fluctuate $\sim e^{-\mathrm{height}/T}$

quantum tunnel $\sim e^{-\sqrt{\mathrm{height}}\times\mathrm{width}/\hbar}$

Encouraging Fact #1

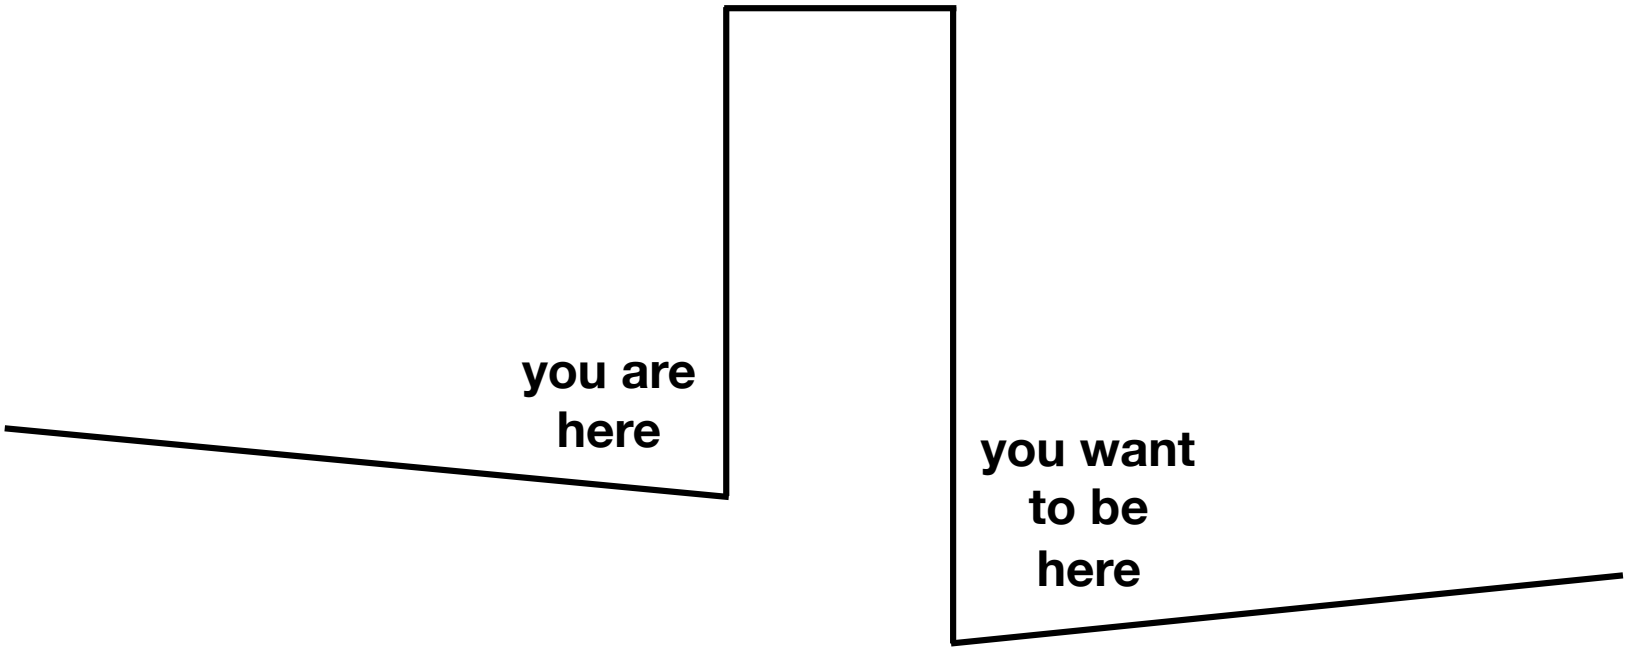

thermally fluctuate $\sim e^{-\text{height}/T}$

quantum tunnel $\sim e^{-\sqrt{\text{height}}\times\text{width}/\hbar}$

quantum beats classical for tall, thin barriers

Encouraging Fact #2

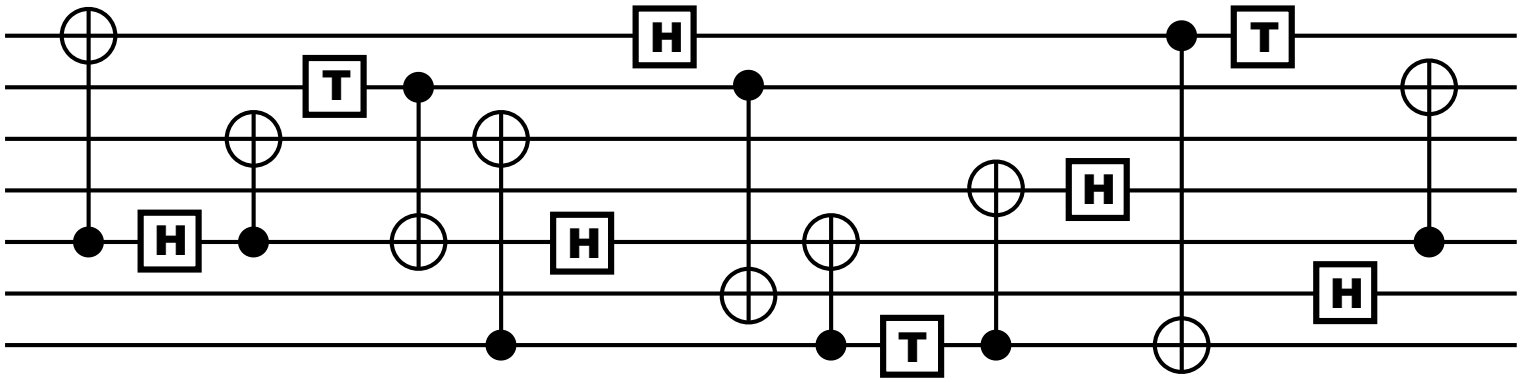

there are short quantum circuits whose output is hard to classically sample (follows directly from BQP not in P)

Encouraging Fact #2

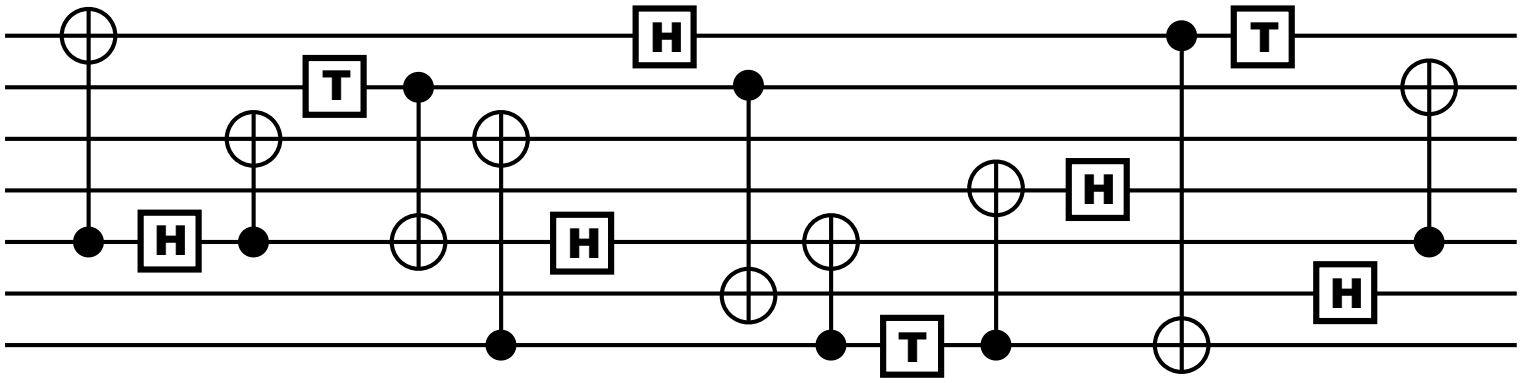

there are short quantum circuits whose output is hard to classically sample (follows directly from BQP not in P)

↓

there are functions that can be expressed with a polynomial size quantum neural network that would require an exponentially large classical neural network

Encouraging Fact #3

Harrow, Hassidim, Lloyd algorithm (2008)

Encouraging Fact #3

Harrow, Hassidim, Lloyd algorithm (2008)

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

Encouraging Fact #3

Harrow, Hassidim, Lloyd algorithm (2008)

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

takes $O(N^c)$ classically

Encouraging Fact #3

Harrow, Hassidim, Lloyd algorithm (2008)

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

takes $O(N^c)$ classically

HHL: with quantum takes $(\log N)^2$!

Encouraging Fact #3

Harrow, Hassidim, Lloyd algorithm (2008)

↓

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

↓

takes $O(N^c)$ classically

HHL: with quantum takes $(\log N)^2$!

↓

subroutine in classical supervised learning is minimize $(\overleftrightarrow{H}\,\vec{x} - \vec{b})^2$

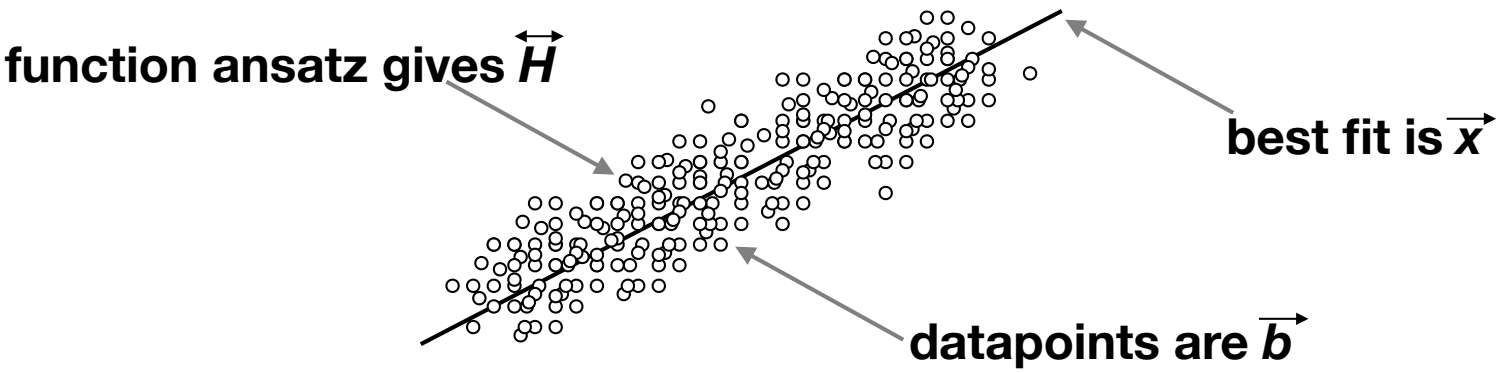

The Power of Quantum Computers (in theory)

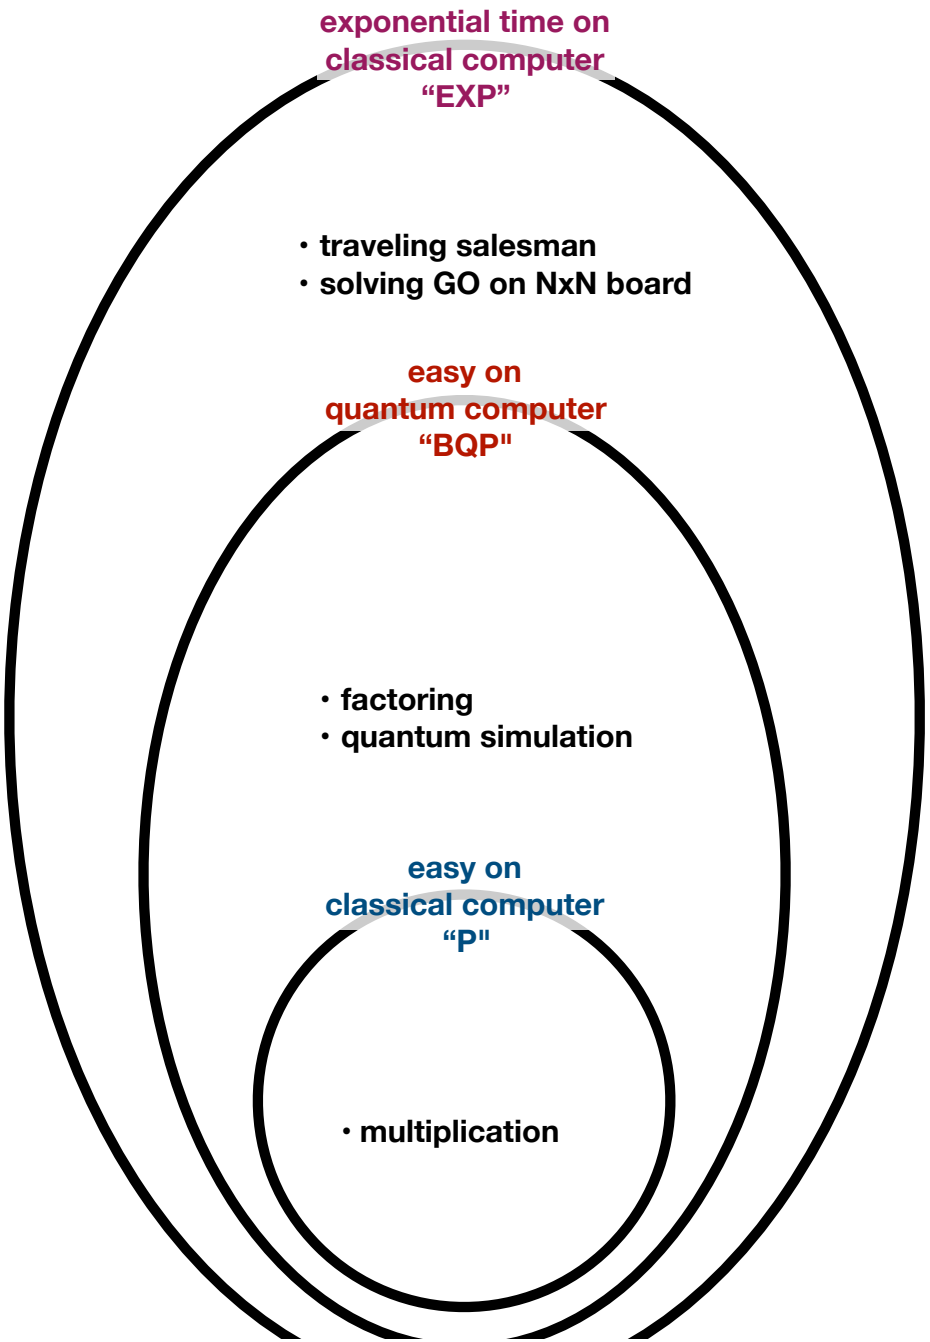

(i) BQP contains P

a quantum computer can (quickly) do
everything a classical computer can (quickly) do

(ii) P does not contain BQP

a quantum computer can do some things quickly
that a classical computer cannot do quickly

↑

ENCOURAGING

(iii) EXP contains BQP

a classical computer can **eventually** do
everything a quantum computer can do

(iv) BQP does not contain EXP

a quantum computer does **not** give
generic exponential speed-up

↑

DISCOURAGING

The Power of Quantum Computers (in practice)

The Power of Quantum Computers (in practice)

- if “you” ever measure, coherence is destroyed
- no cloning principle makes error correction hard
- classical = 10^{-24} errors per gate, quantum = 10^{-2} errors per gate

The Power of Quantum Computers (in practice)

- if “you” ever measure, coherence is destroyed
- no cloning principle makes error correction hard
- classical = 10^{-24} errors per gate, quantum = 10^{-2} errors per gate

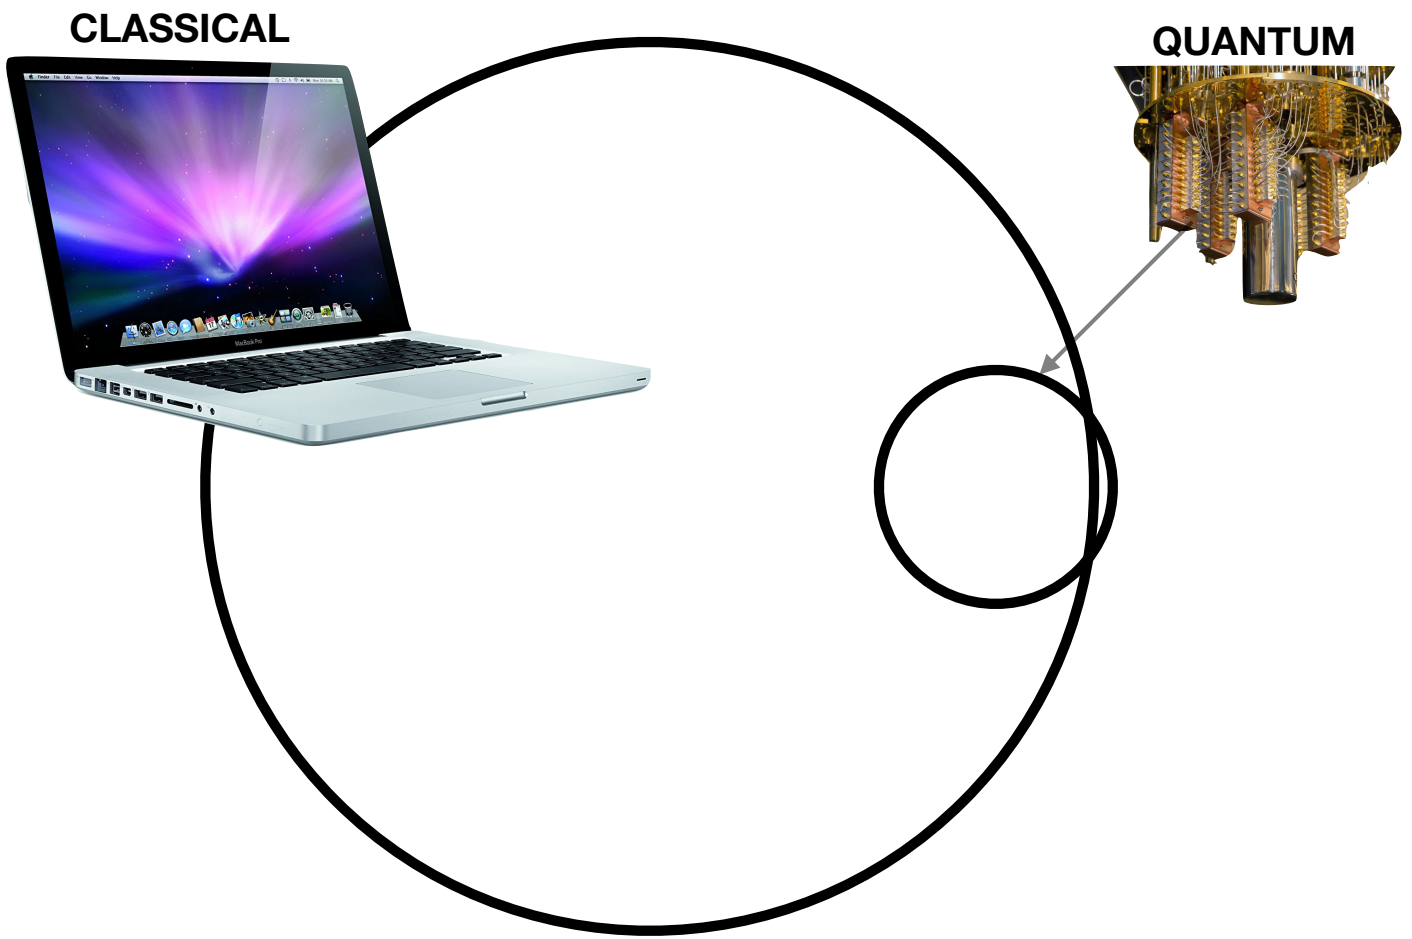

The Power of Quantum Computers (in practice)

- if “you” ever measure, coherence is destroyed
- no cloning principle makes error correction hard
- classical = 10^{-24} errors per gate, quantum = 10^{-2} errors per gate

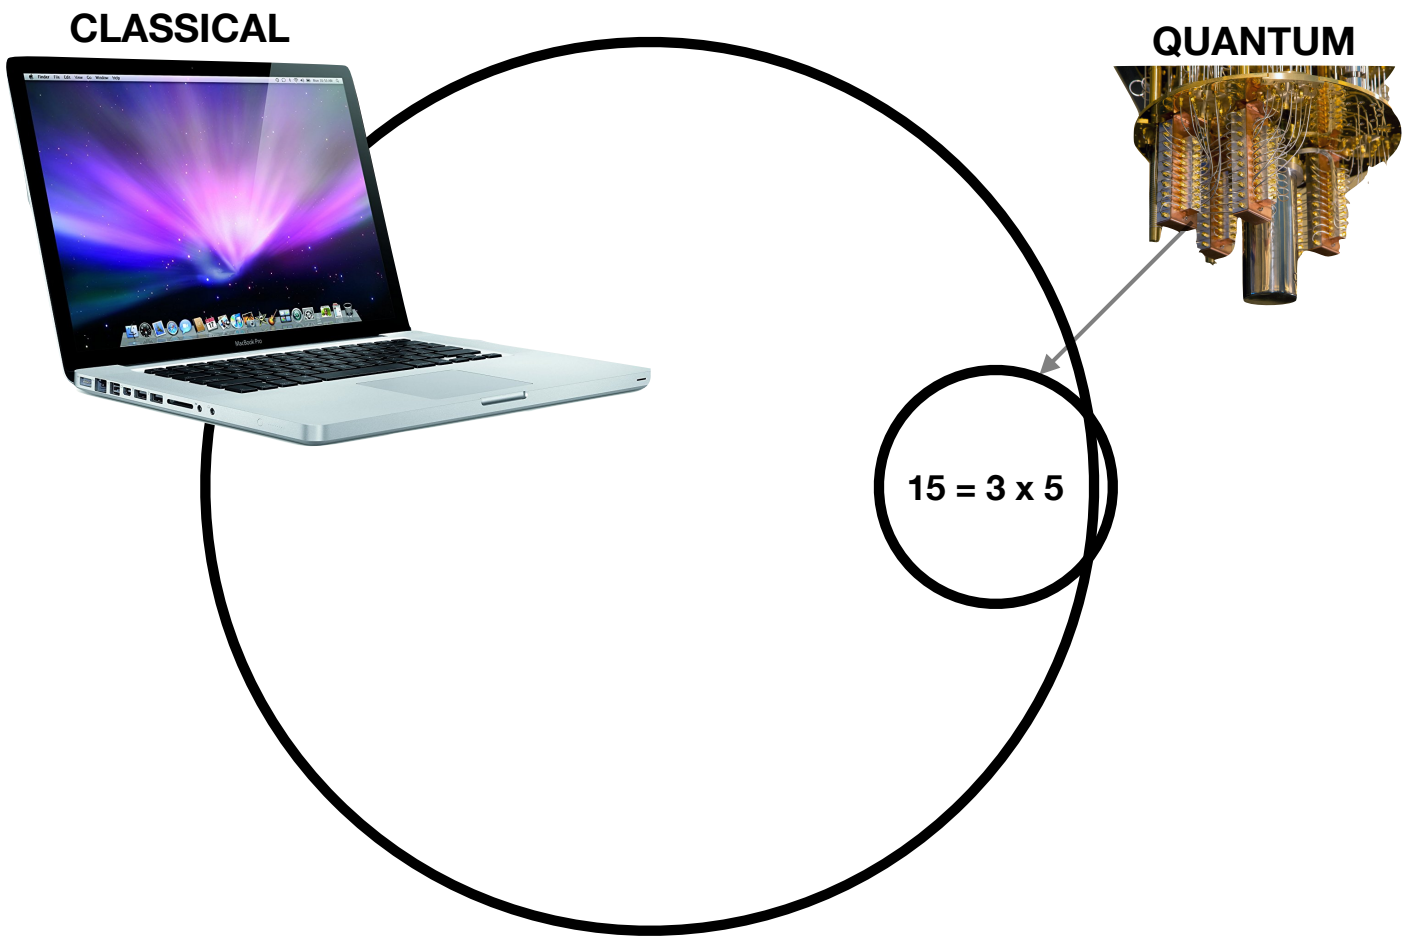

The Power of Quantum Computers (in practice)

- if “you” ever measure, coherence is destroyed
- no cloning principle makes error correction hard
- classical = 10^{-24} errors per gate, quantum = 10^{-2} errors per gate

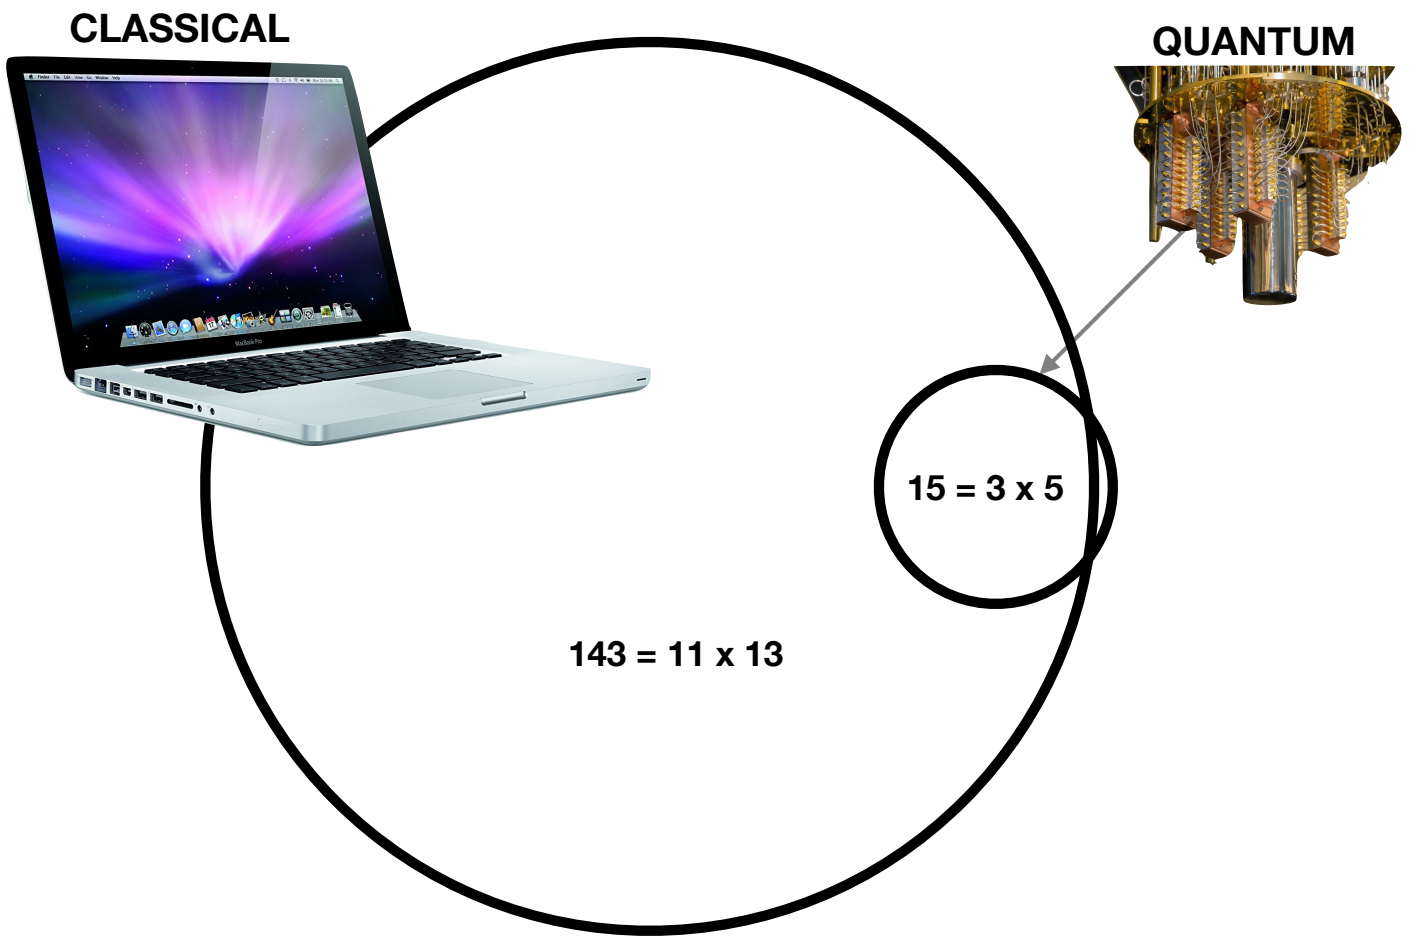

The Power of Quantum Computers (in practice)

- if “you” ever measure, coherence is destroyed
- no cloning principle makes error correction hard
- classical = 10^{-24} errors per gate, quantum = 10^{-2} errors per gate

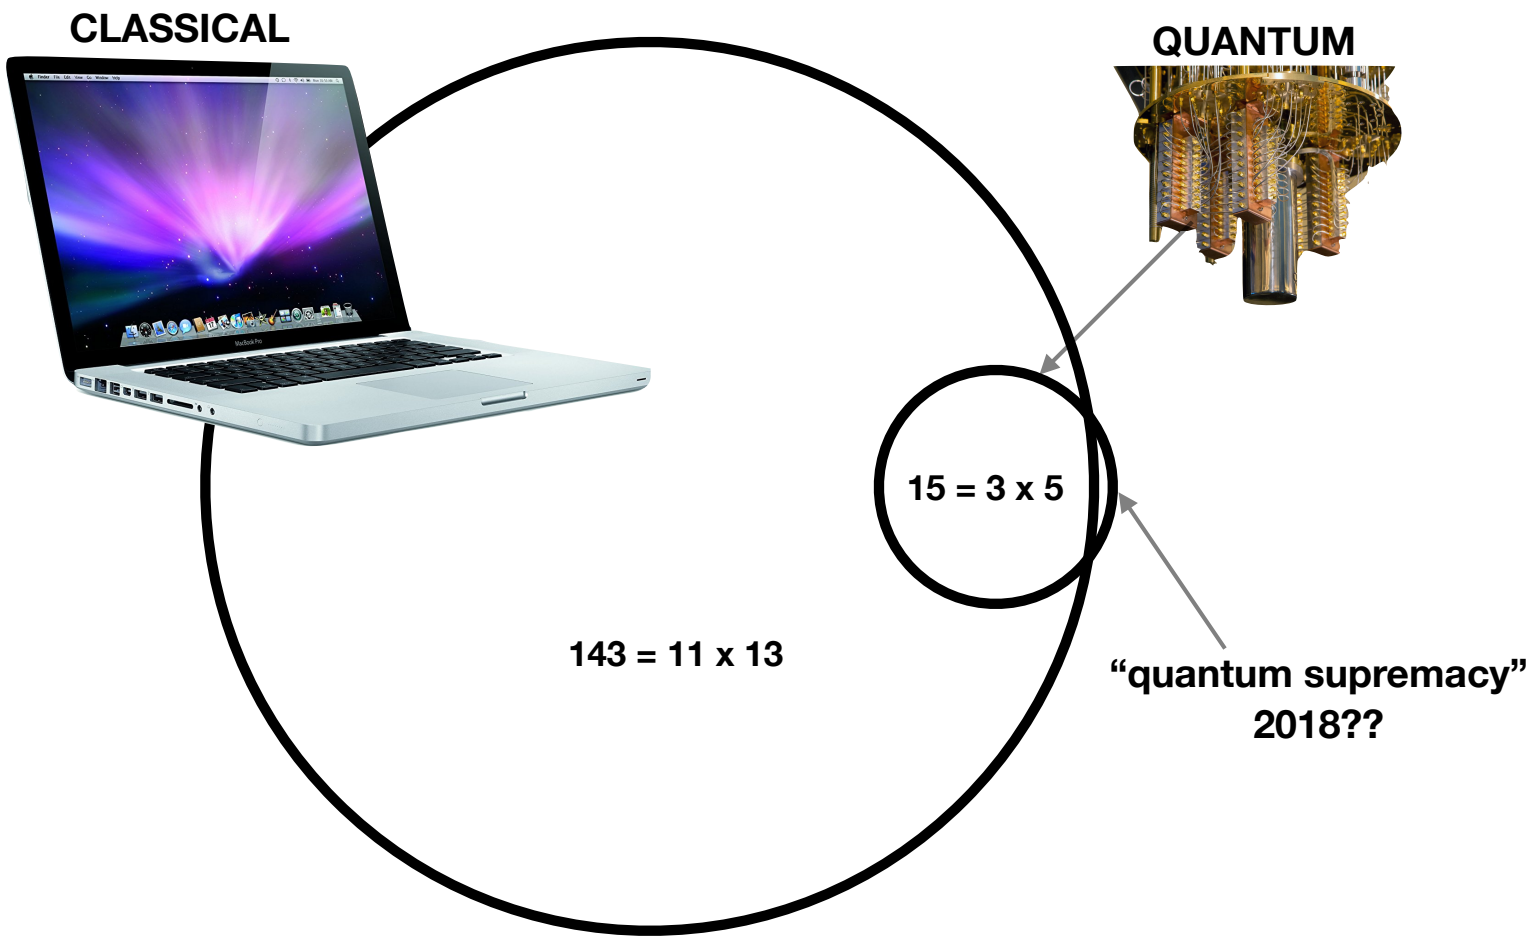

The Power of Quantum Computers (in practice)

“quantum supremacy”
2018??

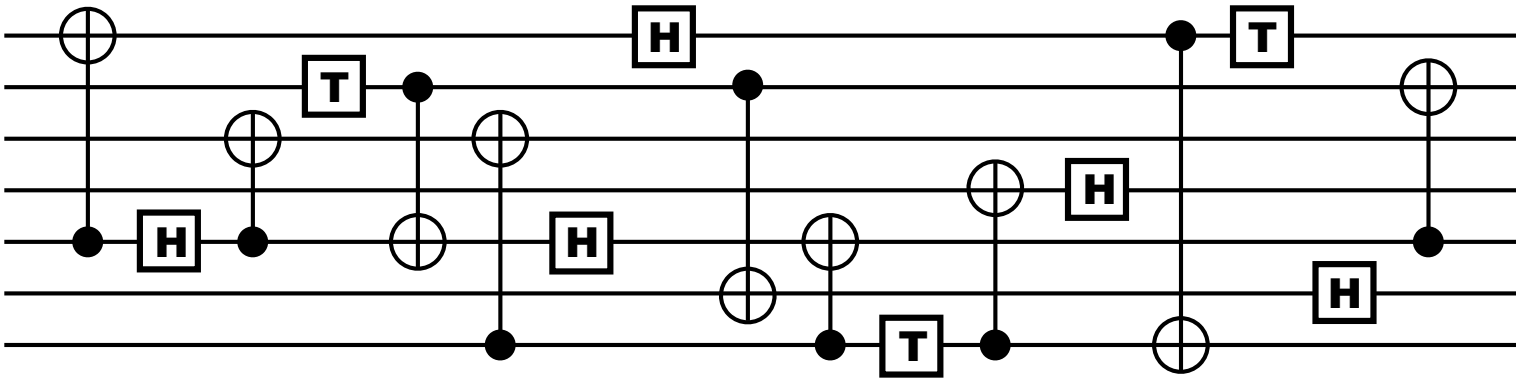

input random state into ~50 qubit random quantum circuit

↓

check on classical supercomputer that you got something consistent with statistical predictions of quantum mechanics

The Power of Quantum Computers (in practice)

“quantum supremacy”
2018??

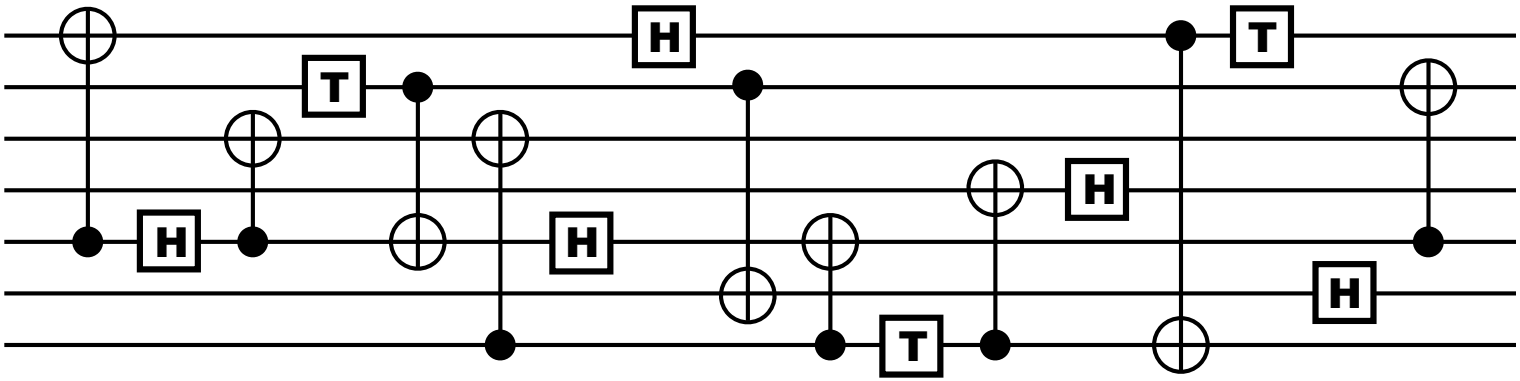

input random state into ~50 qubit random quantum circuit

↓

check on classical supercomputer that you got something consistent with statistical predictions of quantum mechanics

↓

99% noise (error), 1% signal

The Power of Quantum Computers (in practice)

“quantum supremacy”
2018??

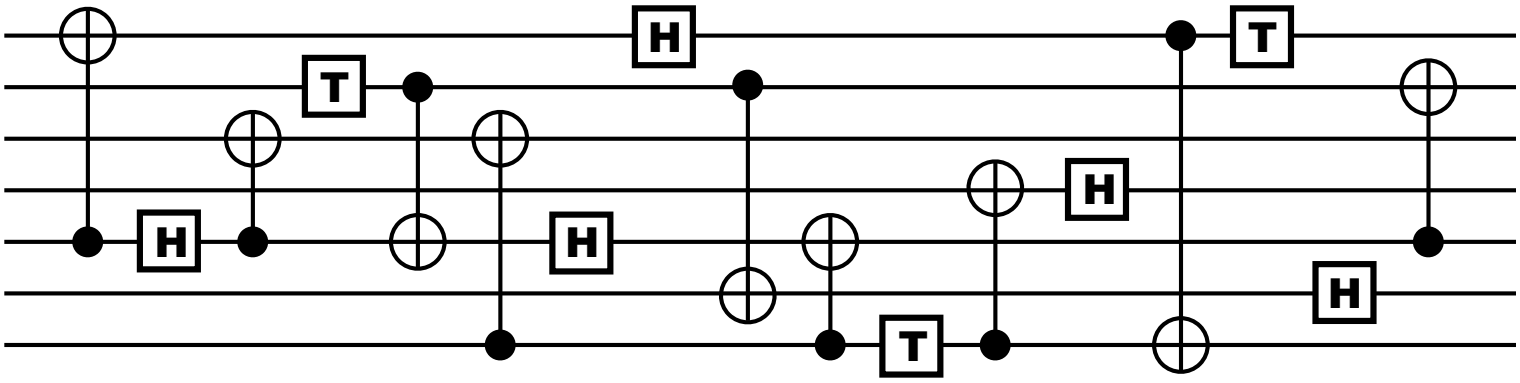

input random state into ~50 qubit random quantum circuit

↓

check on classical supercomputer that you got something consistent with statistical predictions of quantum mechanics

↓

99% noise (error), 1% signal

we’re a long way from cracking RSA!

Example #1: HHL algorithm

given N x N matrix $\overleftrightarrow{H}$

and given vector $\vec{b}$

"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

takes $O(N^c)$ classically

HHL: with quantum takes $(\log N)^2$!

Example #1: HHL algorithm

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

takes $O(N^c)$ classically

HHL: with quantum takes $(\log N)^2$!

Quantum algorithm for solving linear systems of equations
Aram W. Harrow, Avinatan Hassidim, Seth Lloyd
https://arxiv.org/abs/0811.3171

Quantum Machine Learning
Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, Seth Lloyd
https://arxiv.org/abs/1611.09347

Quantum Machine Learning Algorithms: Read the Fine Print
Scott Aaronson
https://scottaaronson.com/papers/qml.pdf

Example #1: HHL algorithm

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

↓

takes $O(N^c)$ classically

HHL: with quantum takes $(\log N)^2$!

$$|b\rangle = \sum_i b_i |i\rangle \longrightarrow |x\rangle = \sum_i x_i |i\rangle$$

b_i ↑ ↓ use $|x\rangle$

Example #1: HHL algorithm

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

takes $O(N^c)$ classically

HHL: with quantum takes $(\log N)^2$!

$$|b\rangle = \sum_i b_i |i\rangle \longrightarrow |x\rangle = \sum_i x_i |i\rangle$$

store *N* bits
in log[*N*] qubits
"QRAM"

b_i

use $|x\rangle$

Example #1: HHL algorithm

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

↓

takes $O(N^c)$ classically

HHL: with quantum takes $(\log N)^2$!

$$|b\rangle = \sum_i b_i |i\rangle \longrightarrow |x\rangle = \sum_i x_i |i\rangle$$

b_i ← store *N* bits in log[*N*] qubits "QRAM"

returns quantum state $|x\rangle$ (**not** vector x_i) → use $|x\rangle$

Example #1: HHL algorithm

given N x N matrix $\overleftrightarrow{H}$
and given vector $\vec{b}$
"find" the vector $\vec{x}$ such that

$$\overleftrightarrow{H}\,\vec{x} = \vec{b}$$

↓

takes O(N^c) classically

HHL: with quantum takes (log N)^2!

$$|b\rangle = \sum_i b_i |i\rangle \xrightarrow[\text{harder than factoring}]{\text{H must be sparse & "well conditioned"}} |x\rangle = \sum_i x_i |i\rangle$$

store N bits in log[N] qubits "QRAM"

b_i

returns quantum state $|x\rangle$ (**not** vector x_i)

use $|x\rangle$

Example #2: Quantum Neural Networks

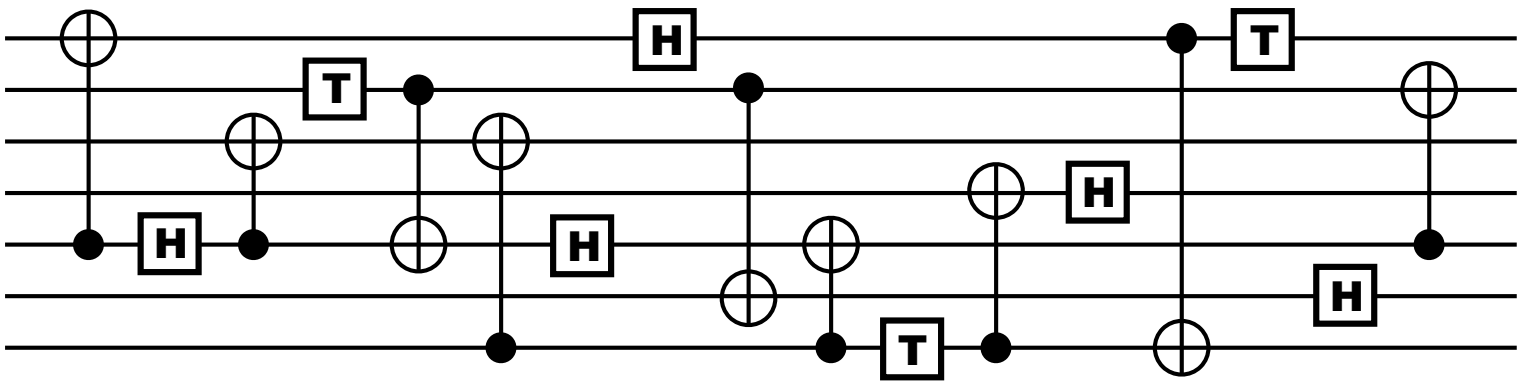

Classification with Quantum Neural Networks on Near Term Processors
Edward Farhi, Hartmut Neven
https://arxiv.org/abs/1802.06002

Barren plateaus in quantum neural network training landscapes
Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, Hartmut Neven
https://arxiv.org/abs/1803.11173

Example #2: Quantum Neural Networks

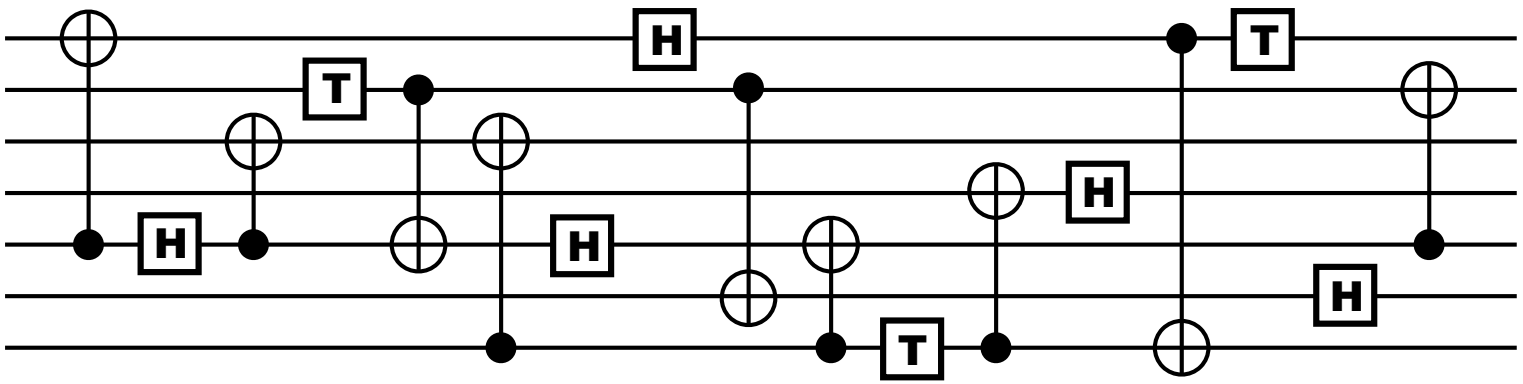

can definitely *represent* some functions that are classically hard

↓

1. **do we *care* about those functions?**
2. **can we *train* it?**

Example #2: Quantum Neural Networks

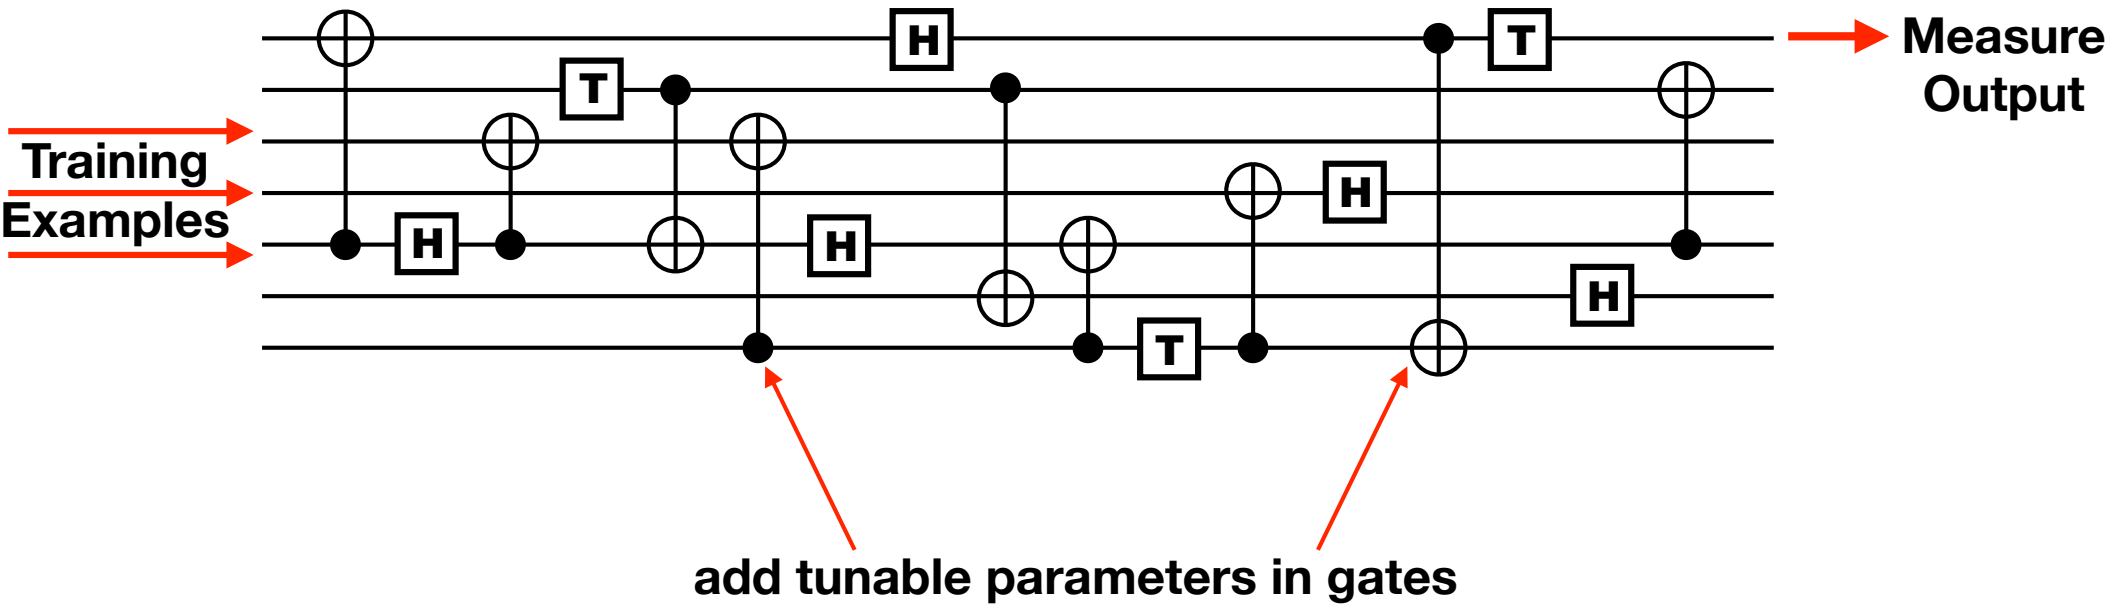

can definitely *represent* some functions that are classically hard

1. do we *care* about those functions?
2. can we *train* it?

Example #3

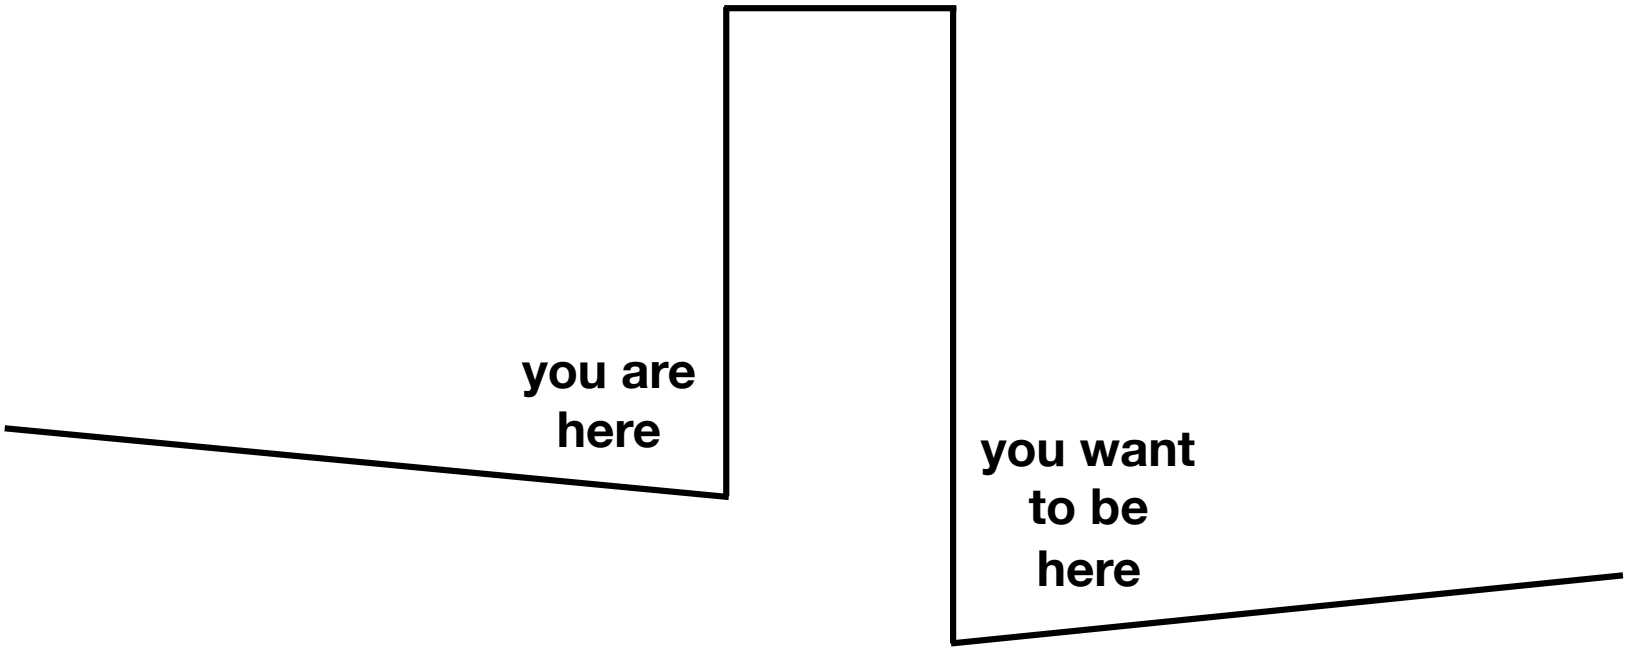

thermally fluctuate $\sim e^{-\mathrm{height}/T}$

quantum tunnel $\sim e^{-\sqrt{\mathrm{height}}\times\mathrm{width}/\hbar}$

quantum beats classical for tall, thin barriers

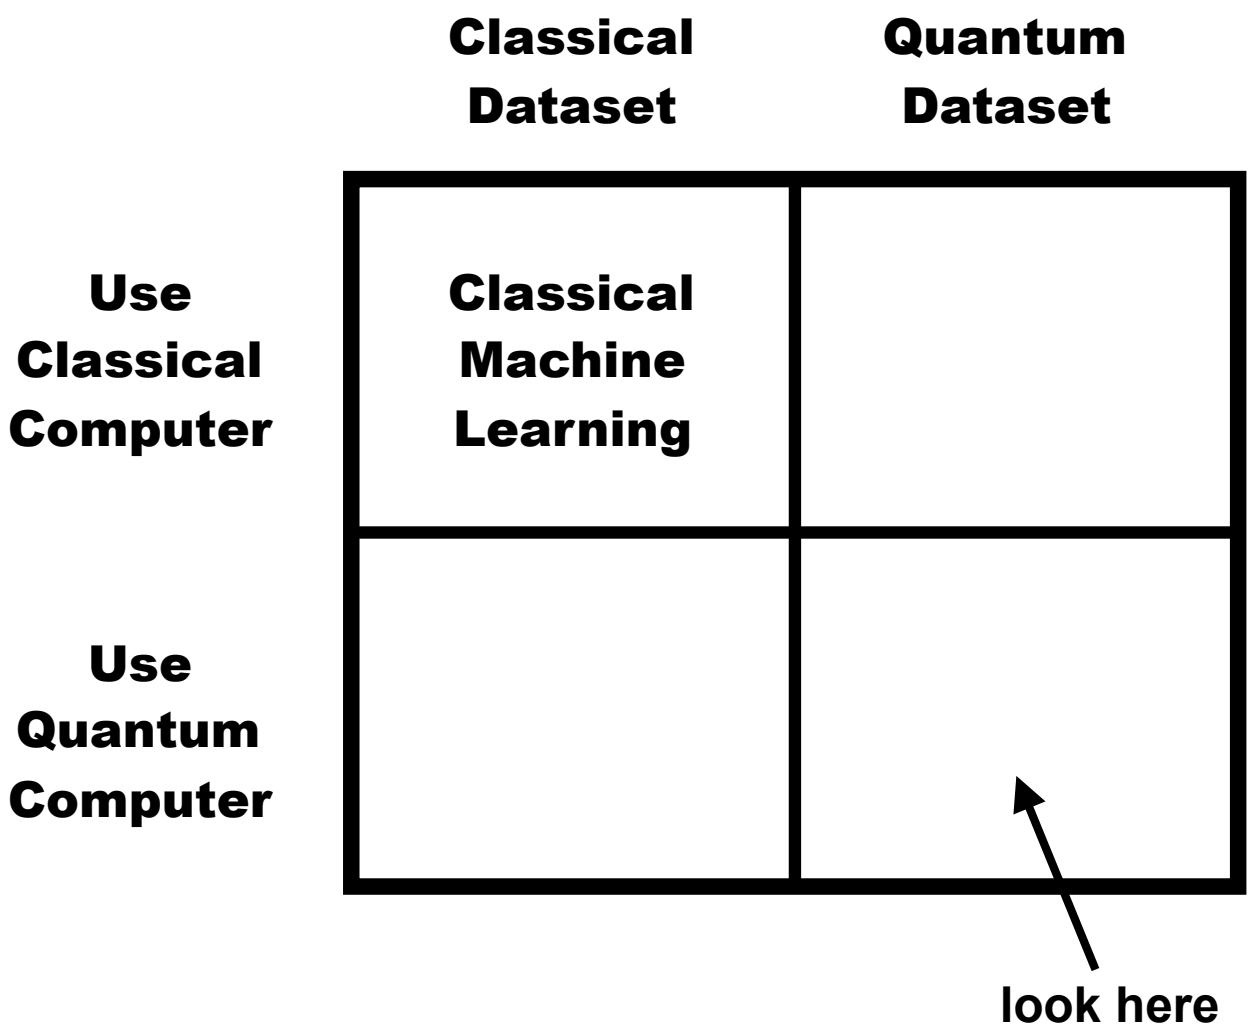
Classical Dataset
Quantum Dataset
Use Classical Computer
Use Quantum Computer
Classical Machine Learning
look here