Quantum learning algorithms imply circuit lower bounds

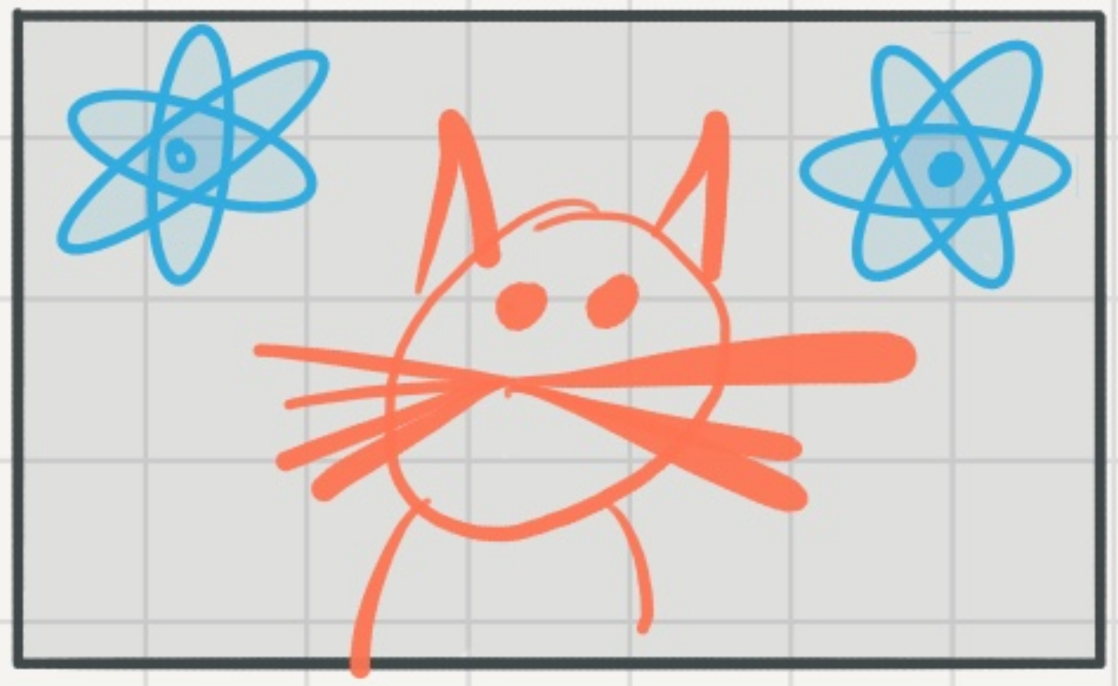

$|1\rangle$: cat

$|0\rangle$: Not cat

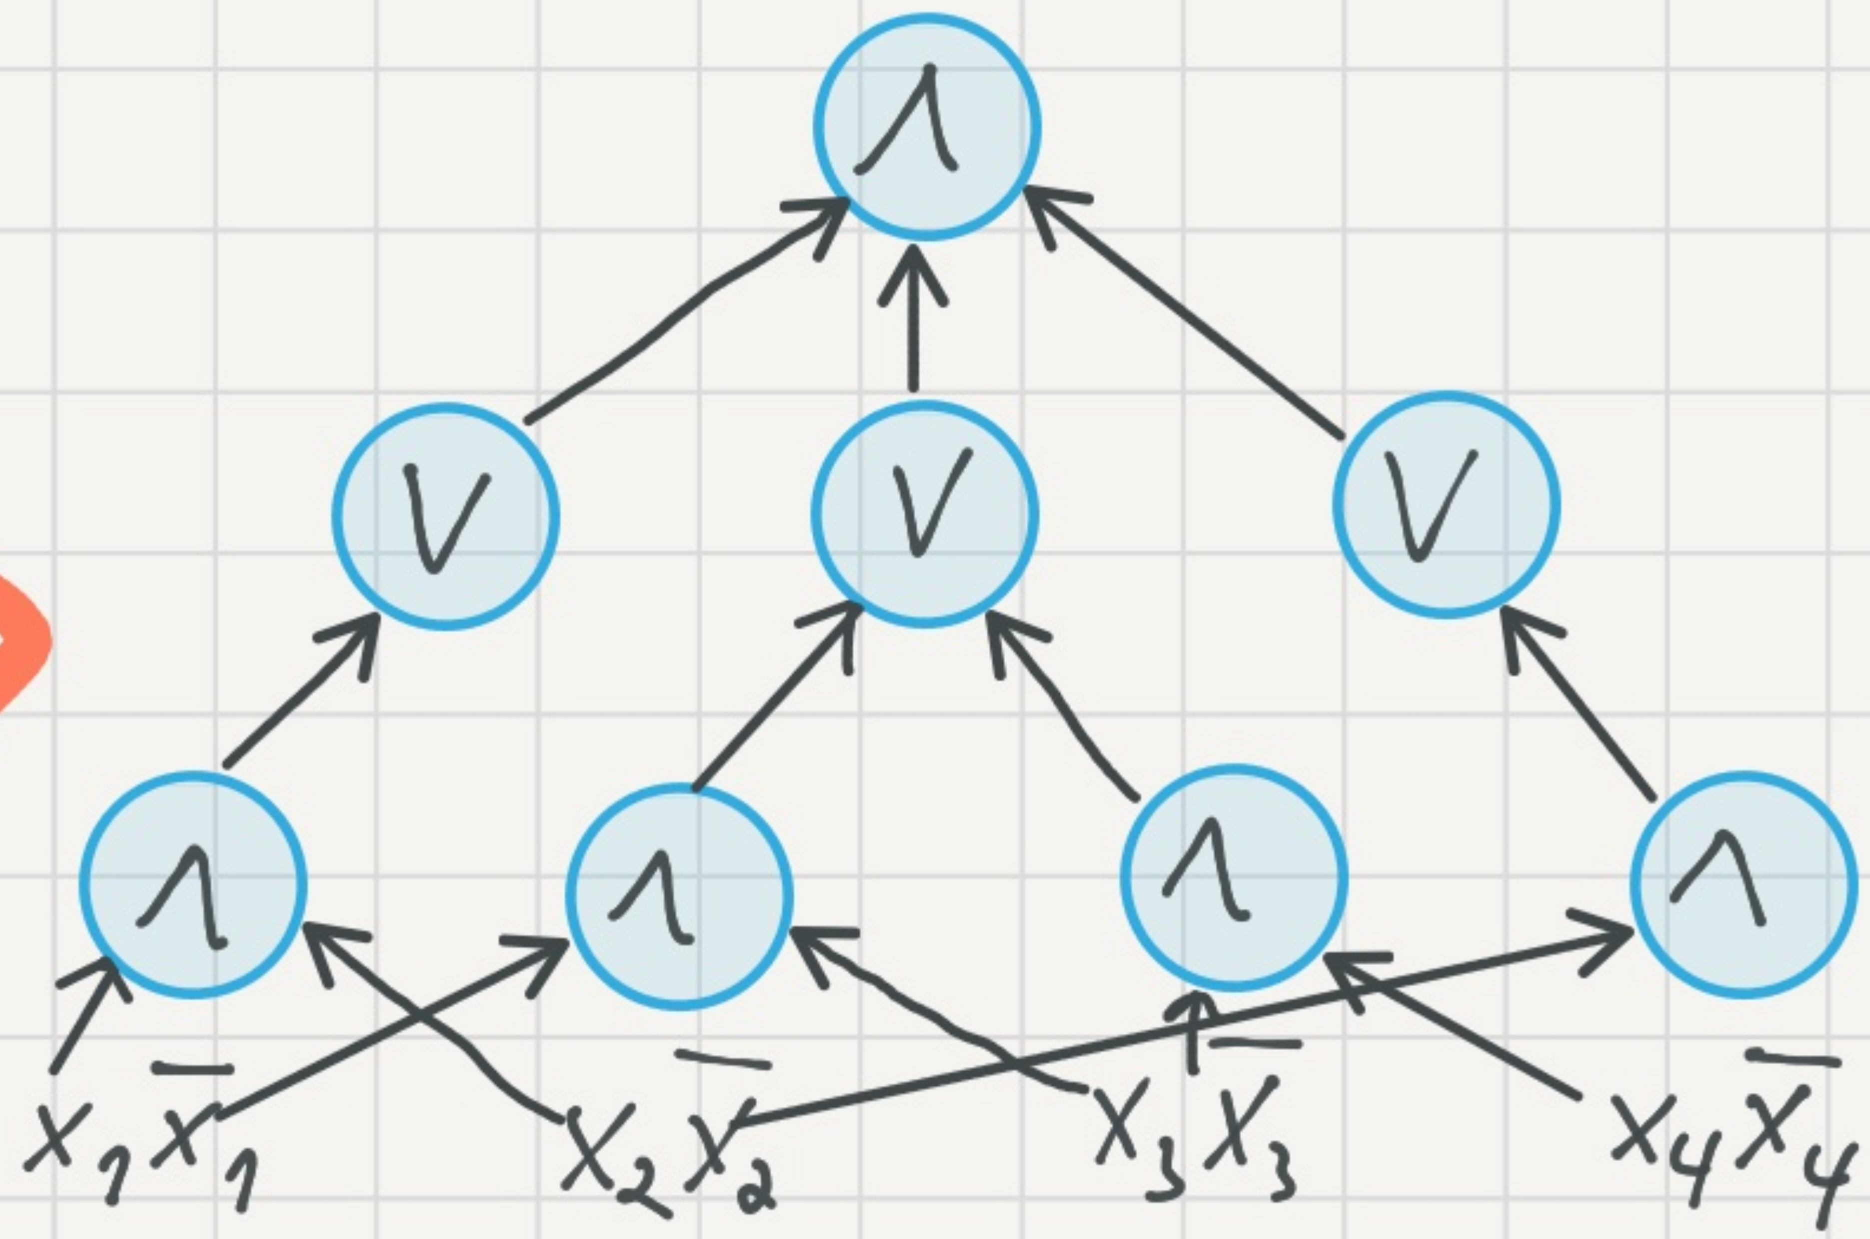

Arunachalam – Grilo – G – Oliveira – Sundaram

IBM Quantum, CNRS, Warwick, Warwick, Microsoft Research


"Achieve quantum Speedups"
Quantum Algorithms
Learning Theory
"Design fast learning algorithms"
Complexity Theory
"Understand the limits of algorithms"
This work

The take-home message

Any marginal improvement on "trivial" quantum learners will imply breakthrough circuit lower bounds

OPTIMIST

"A new path to lower bounds using the power of quantum information!"

PESSIMIST

"An explaination why quantum learners are so hard to design!"

Circuit lower bounds

A million dollar question: **P** vs **NP**

circuit class

$\mathcal{C}[s(n)]$

size

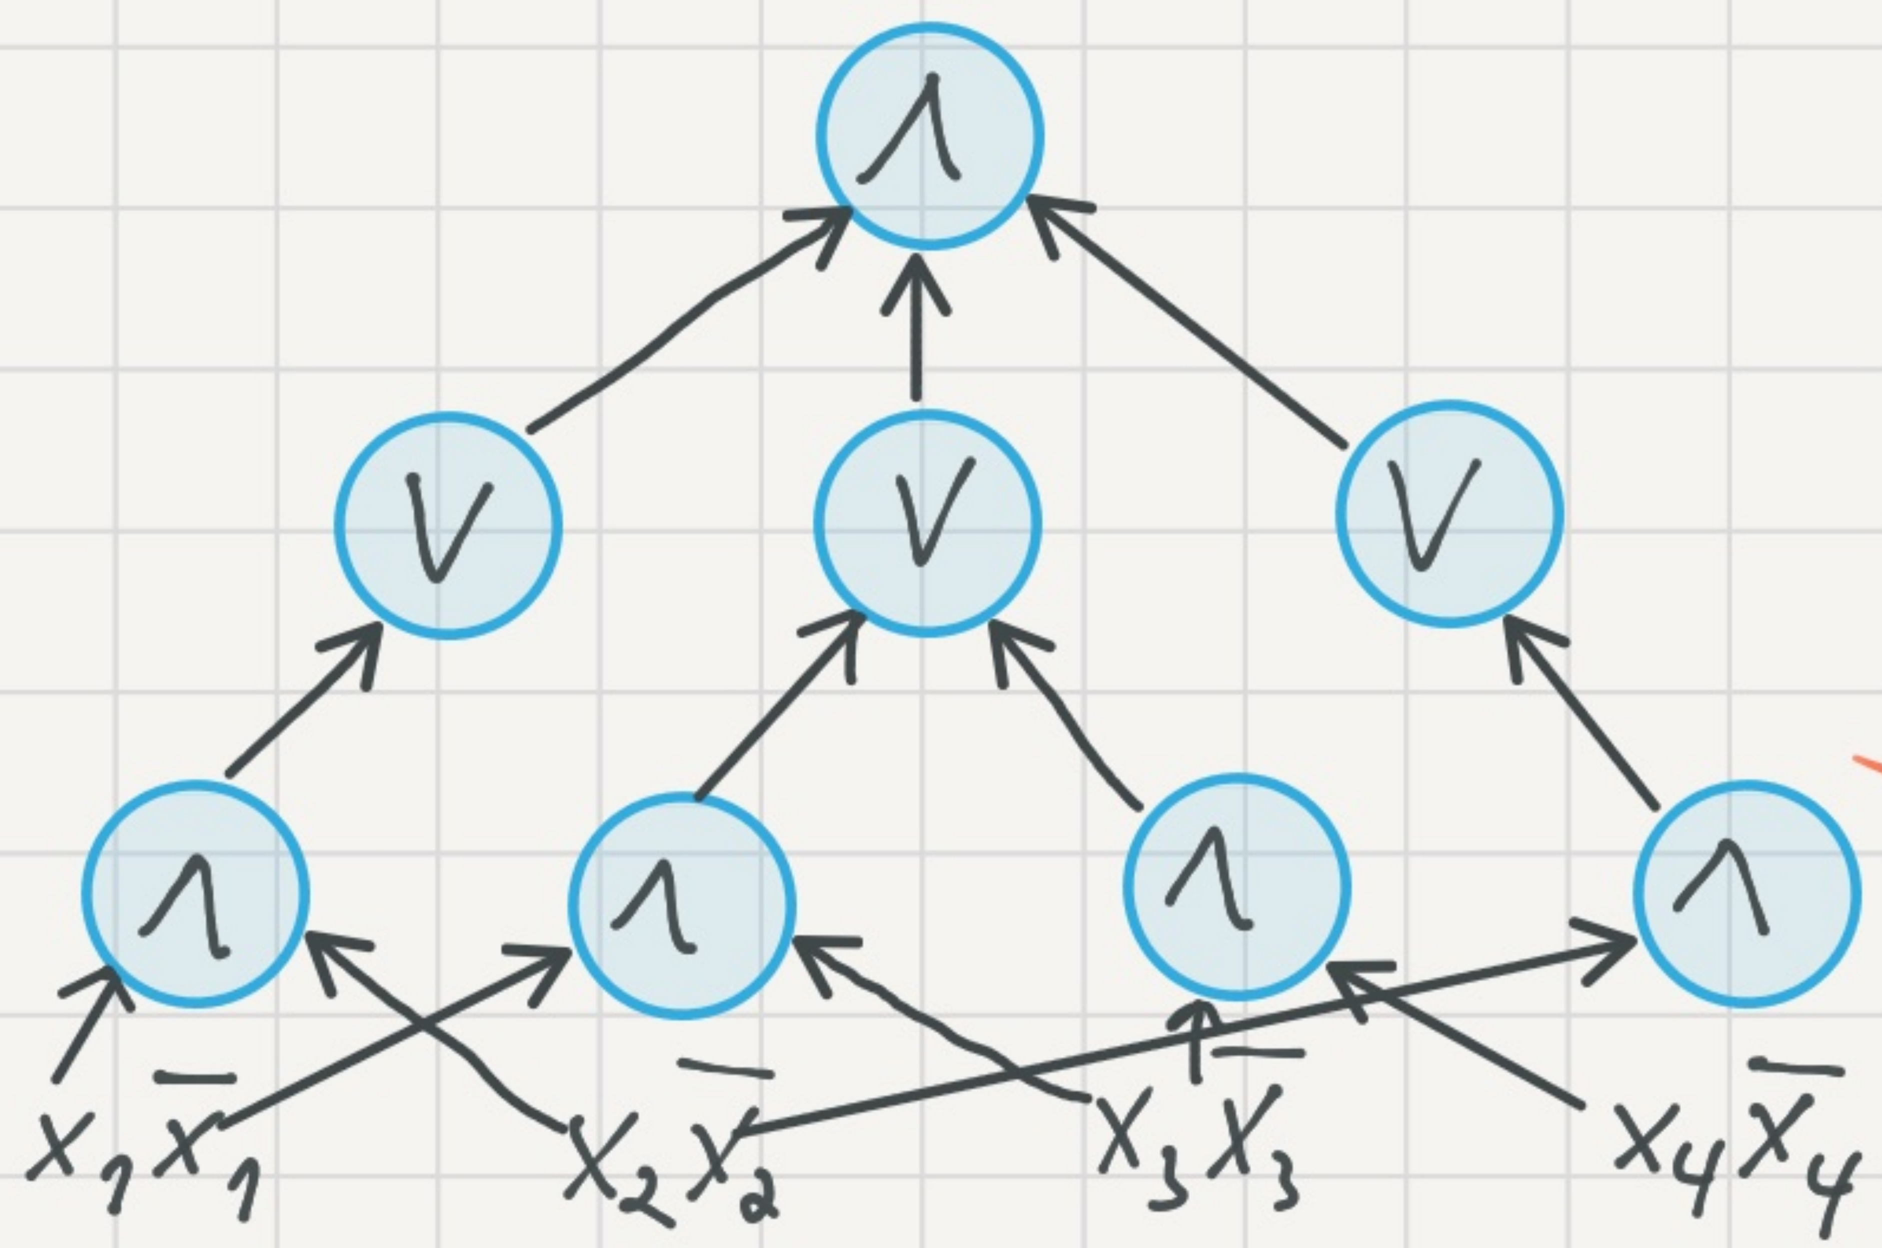

Approach dating back to Shannon, 1949

Turing machine are complex,

but circuits are simple, combinatorial objects...

Back in the 80's:

"It's a new day, it's a new dawn, and I'm feeling good!"

— Nina Simone

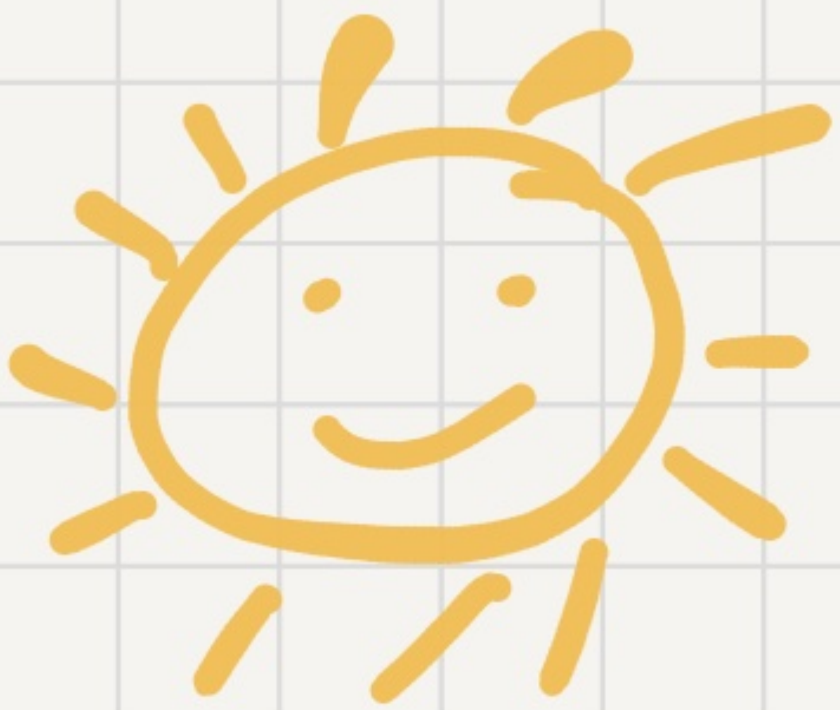

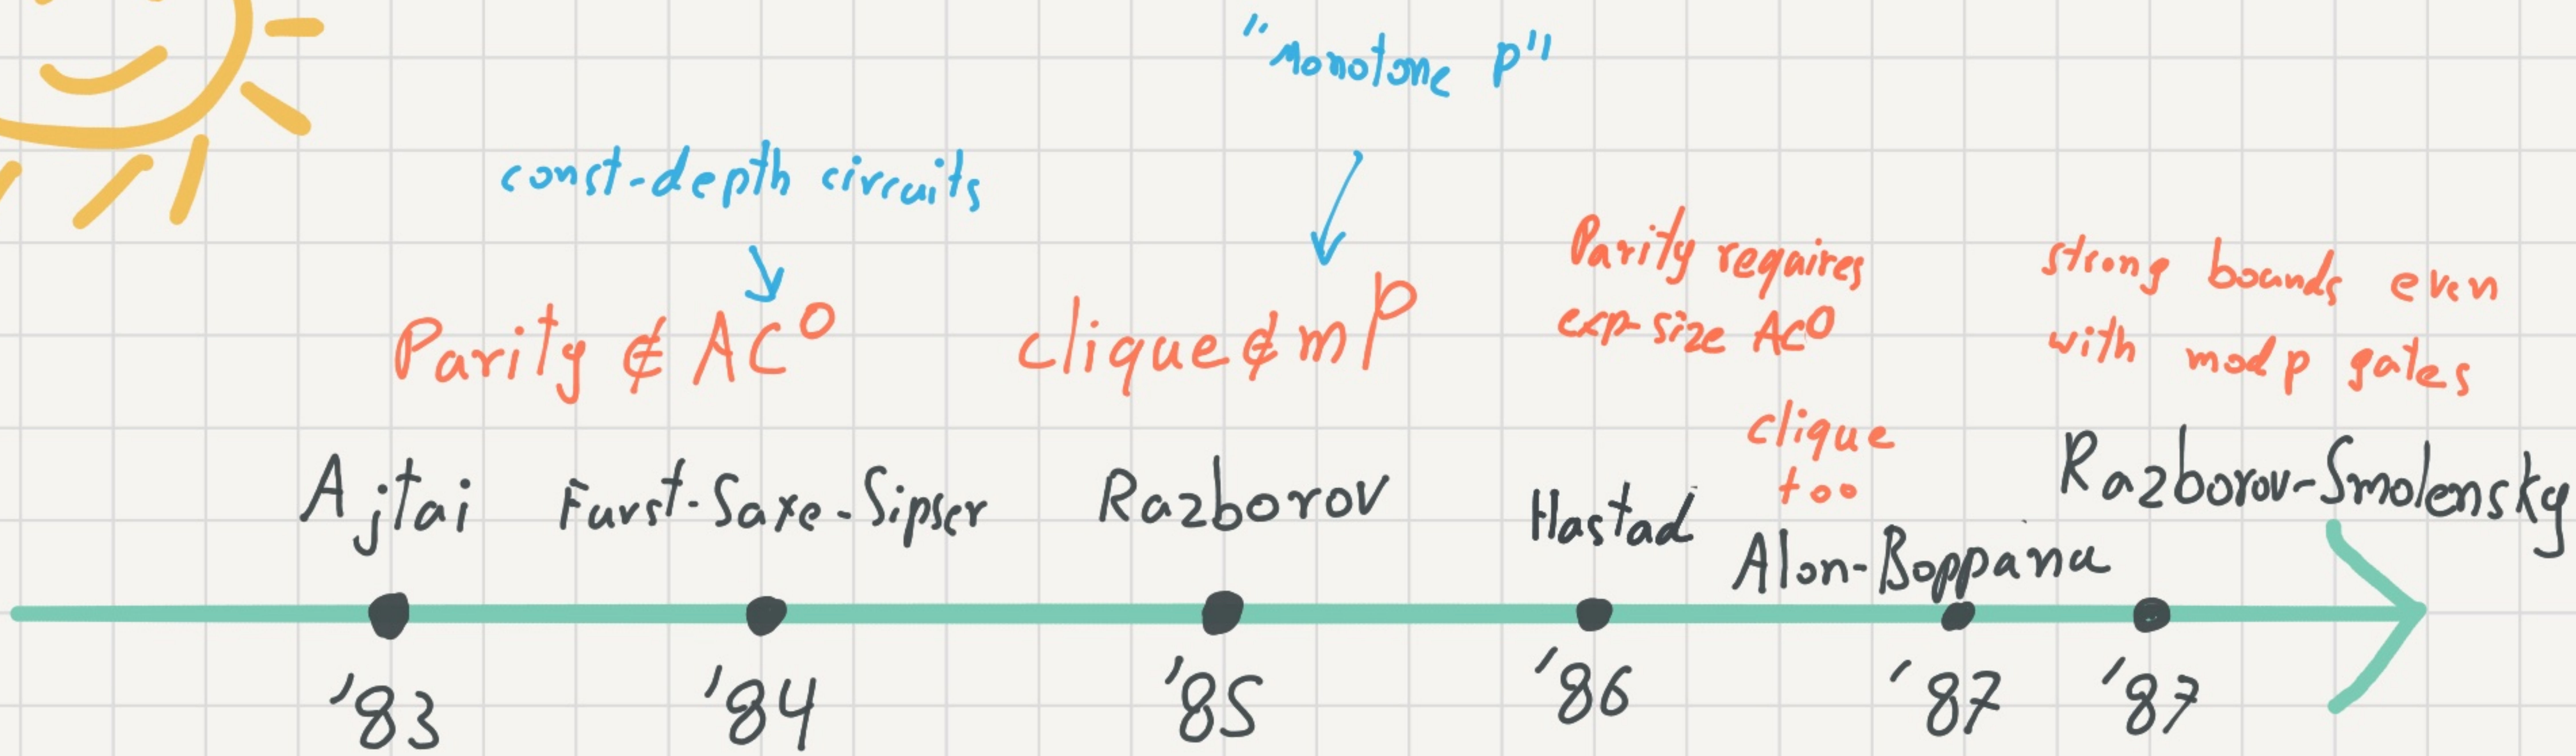

A meteoric rise for circuit lower bounds!

Surely, P vs NP is on the horizon...

Since the 90's:

"I find it hard to tell you,
I find it hard to take. When people run in circles it's a very, very mad world"
— Gary Jules

From here, things went (mostly) downhill...

Belief: $NP \not\subseteq P/Poly$ (o/w, complexity doomsday: PH collapses by Karp-Lipton)

poly-size circuits

What we actually know: NEXP could have poly-size circuits...

BQE could have depth-2 threshold circuits...

A new paradigm

Circuit lower bounds are hard to prove!

A major breakthrough: $\mathsf{NEXP} \not\subseteq \mathsf{ACC}^0$ [Williams '11]

const-depth
$\vee, \wedge, \neg,$ MOD

Big idea: derive lower bounds from algorithms!

Quantum learning Theory

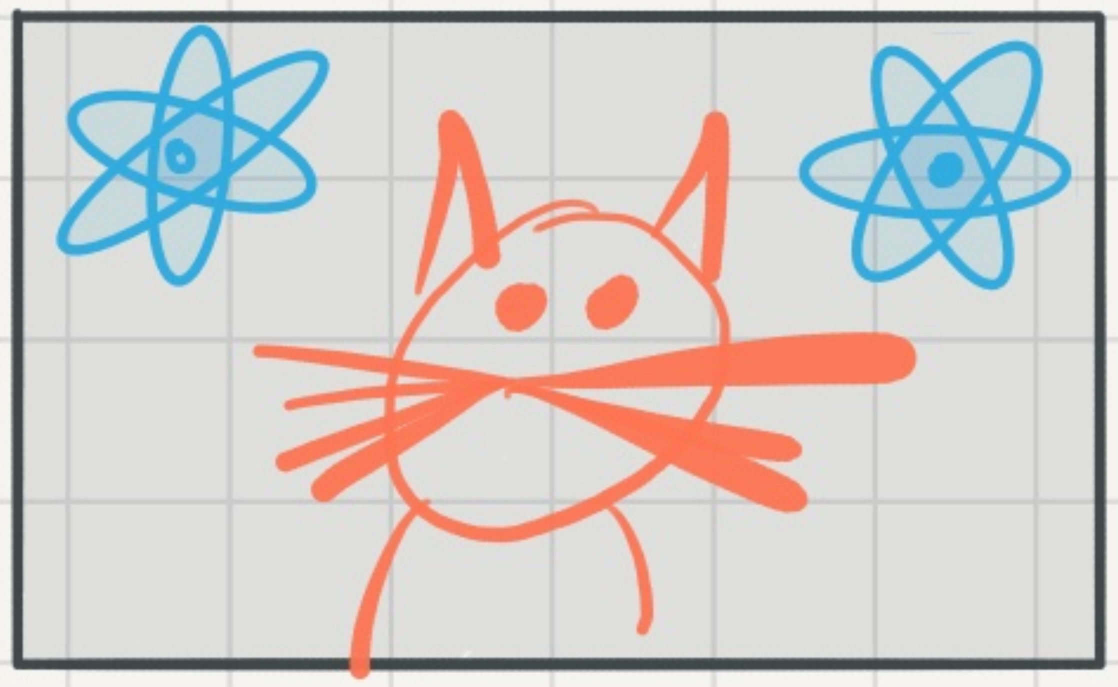

$|1\rangle$: cat

$\rightarrow$ $|0\rangle$: Not cat

Setting: a known class $\mathcal{C}$

an unknown function $f \in \mathcal{C}$

query access to f (on input $|x\rangle|b\rangle$, return $|x\rangle|b\oplus x\rangle$)

Goal: w.p. $1-\delta$, output a circuit c s.t.

$$\Pr_x[c(x) \neq f(x)] \leq \varepsilon$$

For our result:
the stronger the model,
the better!

Quantum learning Theory

Setting: a known class $\mathcal{C}$

an unknown function $f \in \mathcal{C}$

query access to f (on input $|x\rangle|b\rangle$, return $|x\rangle|b \oplus x\rangle$)

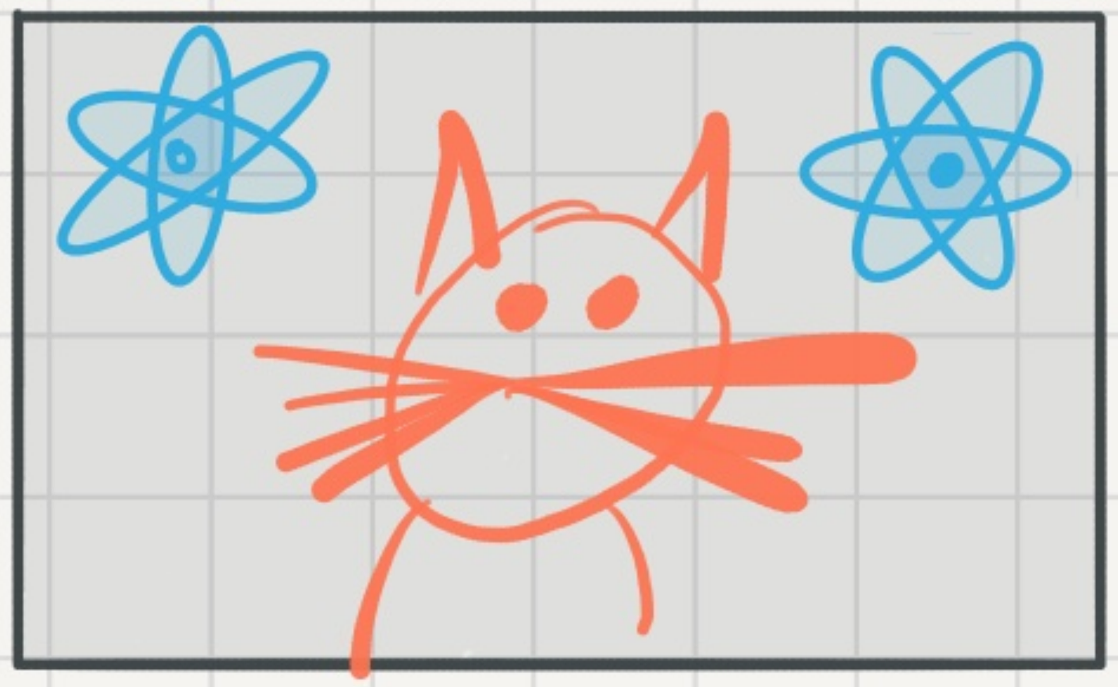

$\rightarrow$ $|1\rangle$: cat

$|0\rangle$: Not cat

More generally: w.p. $1-\delta$, output a quantum circuit U s.t.

$$\mathbb{E}_x\left[\|\Pi_{f(x)} U|x\rangle|0^a\rangle\|^2\right] \leq \varepsilon$$

$\Pi_{f(x)} = |f(x)\rangle\langle f(x)| \otimes I$

Can quantum queries help?

So far, little is known...

Formula-SAT $\in$ BQTIME$[2^{n/2}]$

e.g.

Many classical negative results no longer hold!

Quantum query algorithms admit many speedups

Foundamental question

Are any quantum learning speedups possible?

e.g.

learn TC_0 circuits in quantum time $2^{n/2}$?

Our result

Theorem

If a class $\mathcal{C}$ of poly-size concepts can be learned with error $\varepsilon \leq \frac{1}{2} - \gamma$ in quantum-time $o(\gamma^2 \cdot \frac{2^n}{n})$, then $\mathsf{BQE} \not\subseteq \mathcal{C}$.

First general connection between quantum algorithms & complexity lower bounds

Our result

Theorem

If a class $\mathcal{C}$ of poly-size concepts can be learned with error $\varepsilon \leq \frac{1}{2} - \gamma$ in quantum-time $o(\gamma^2 \cdot \frac{2^n}{n})$, then $\mathrm{BQE} \not\subseteq \mathcal{C}$.

Example: If poly-size depth-2 threshold circuits can be learned in quantum-time $o(2^n/n)$

=> new circuit lower bound!

"Minor" algorithmic progess => major progress in complexity theory!

Tightness of the result

Consider 2 "trivial" quantum learners:

1) Query everything

Time: 2^n, Error: 0

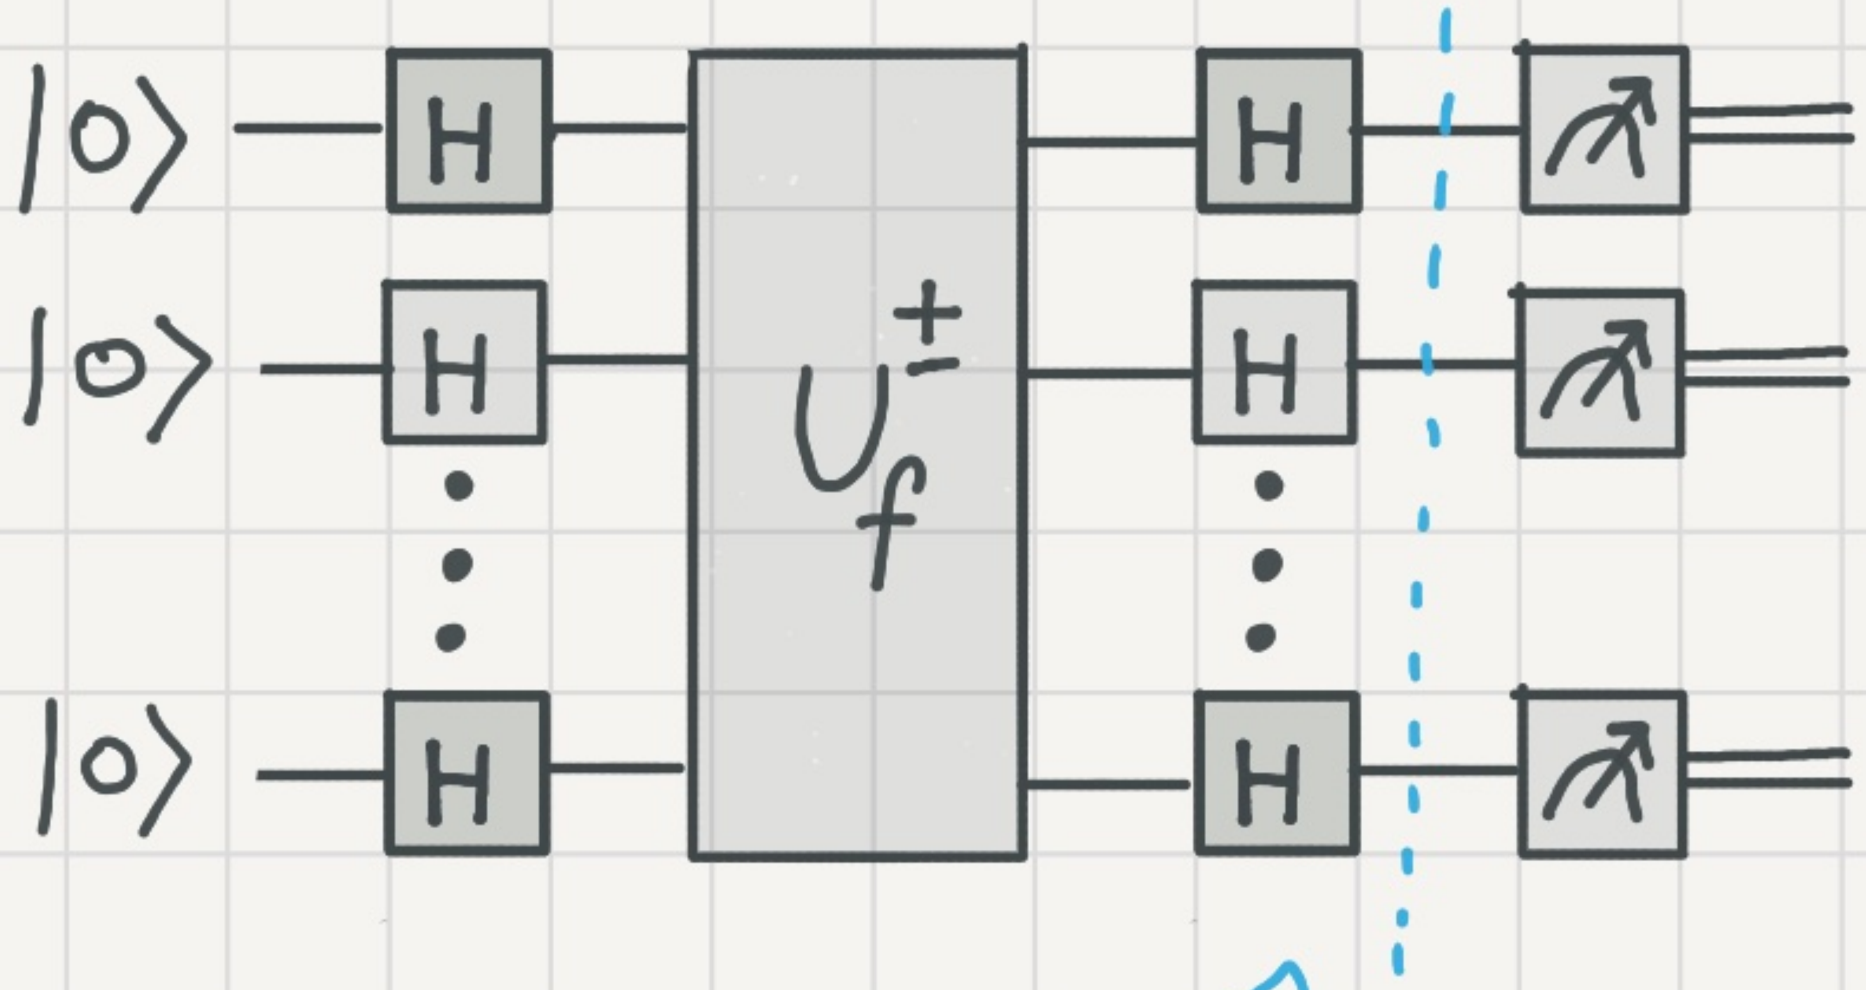

$|\psi\rangle = \sum_s \hat{f}(s)|s\rangle$

$\hat{f}(s) = \langle \chi_s, f\rangle$

2) Fourier sampling

Time: poly(n), Error: $\frac{1}{2} - \Omega(2^{-n/2})$

Our result: Non-trivial quantum learners => circuit lower bounds

Proof Overview

1) Quantum learners => $\underbrace{\text{quantum natural-properties.}}$

$\mathcal{C}$ admits a quantum algorithm that:

Structure vs randomness ← I) rejects $f \in \mathcal{C}$

II) accepts a dense set

III) runs in time $\mathrm{poly}(|f|) = 2^{O(n)}$

Idea: If we can q-learn $\mathcal{C}$, we can distinguish $f \in \mathcal{C}$ from a random f

Proof Overview

1) Quantum learners => quantum natural-properties.

2) If $\mathsf{PSPACE} \not\subseteq \mathsf{BQTIME}[2^{n^\alpha}]$, there exists a PRG fooling uniform quantum circuits.

First conditional pseudorandom generator secure against uniform quantum computations

(Quantizing hardness amplification results)

Proof Overview

1) Quantum learners => quantum natural-properties.

2) If PSPACE $\not\subseteq$ BQTIME$[2^{n^\alpha}]$, there exists a PRG fooling uniform quantum circuits.

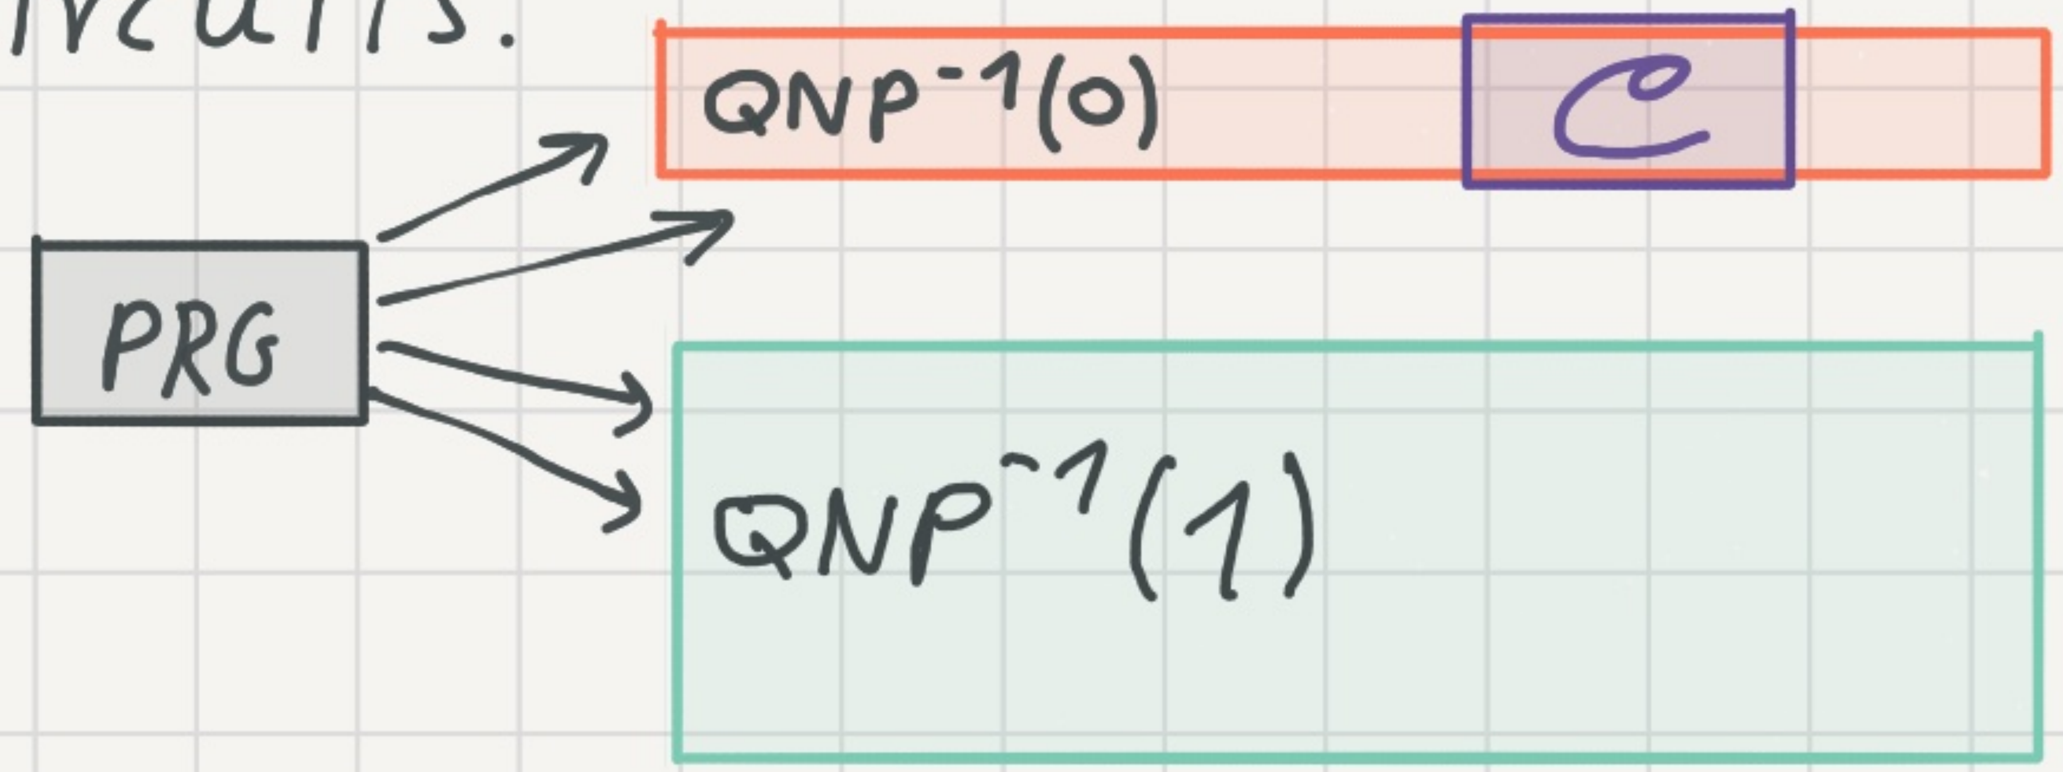

3) "Win-win" argument:

- If PSPACE $\subseteq$ BQTIME$[2^{n^\alpha}]$, lower bounds via diagonalization.
- If PSPACE $\not\subseteq$ BQTIME$[2^{n^\alpha}]$, use PRG to fool the quantum natural-property to hit a hard function.

Open problems

1) Design non-trivial quantum learners.

2 birds, 1 stone: new learners, new lower bounds!

2) Quantum natural properties against ACC^0?

Torus polinomials? Symmetric rank?

3) Stronger PRGs against uniform quantum circuits?

Thank you!