QUANTONATION

Crushing the Barrier to Entry: How Agentic AI is Opening Up Quantum Application Development

Christophe Jurczak

Quantonation Investor Day · July 2, 2026

THE BARRIER TO ENTRY TO PROGRAM A QPU IS REAL

pulser_z3.py

```
from pulser import Register, Sequence, Pulse
from pulser.waveforms import RampWaveform

reg = Register.triangular_lattice(rows=5,
        cols=5, spacing=5.0)
seq = Sequence(reg, AnalogDevice)
seq.declare_channel("rydberg","rydberg_global")

Omega    = 2*np.pi*3.3          # Rabi freq
d0, df   = -6*Omega, 6*Omega    # detuning
seq.add(Pulse.ConstantDetuning(
  RampWaveform(4000,0,Omega), d0, 0), "rydberg")

# blockade radius? register embedding?
# noise model? SPAM? bond dimension?
# adiabatic ramp time? ...
```

…and that's the easy part.

Register geometry. Pulse sequences. A Hamiltonian. Noise models. Bond dimensions. SPAM correction.

Conclusion:

first, go get a PhD.

AGENTIC AI: ONE INVESTIGATOR, END TO END

Agentic tooling now lets a single person carry an application from idea to result — the way it is already happening across every other industry.

I tested it on a hobby — my collection of Renaissance books.

- The 1544 print is dense, heavily-accented early Italian. Standard OCR is unusable; even Transkribus HTR misses ~16% and drops a quarter of the accents. So I fine-tuned a vision-language model — 2.48%.
- Shipped the model and the data publicly — Zenodo and Hugging Face.

I trained and released an ML model as a side project. THAT is the barrier collapsing ☺

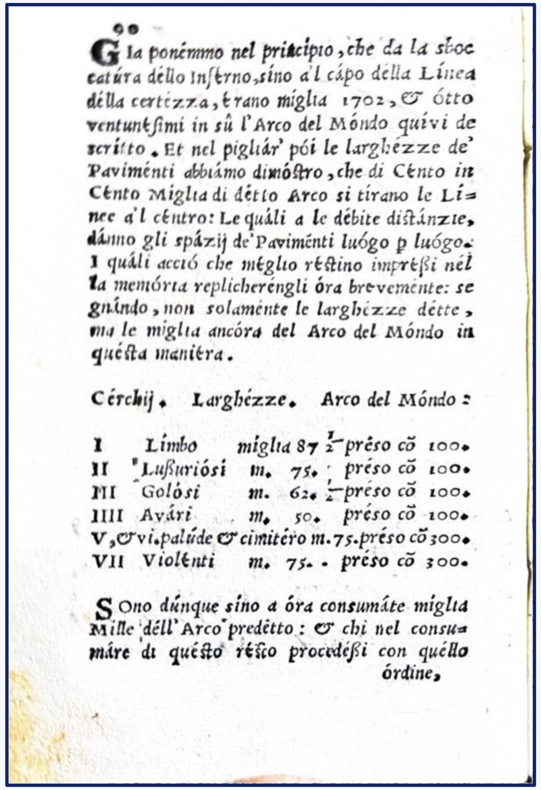

90

Gla ponémmo nel principio, che da la sboccatúra dello Infrerno, sino a'l cápo della Línea della certézza, erano miglia 1702, & ótto ventunésimi in sû l'Arco del Móndo quivi descritto. Et nel pigliár' pói le larghézze de' Paviménti abbiámo dimóstro, che di Cénto in Cénto Miglia di détto Arco si tirano le Línee a'l céntro: Le quáli a le débite distánzie, dánno gli spázij de' Paviménti luógo p luógo. I quáli acció che mèglio restino impréssi nélla memória replicheréngli óra brevemente: seguéndo, non solaménte le larghézze détte, ma le miglia ancóra del Arco del Móndo in quésta manièra.

Cérchij.	Larghézze.	Arco del Móndo:
I Límbo	miglia 87 ½	préso cõ 100.
II Lussuriósi	m. 75.	préso cõ 100.
III Golósi	m. 62. ½	préso cõ 100.
IIII Avári	m. 50.	préso cõ 100.
V, & vi. palúde & cimitéro	m. 75.	préso cõ 300.
VII Violénti	m. 75.	préso cõ 300.

SOno dúnque sino a óra consumáte miglia Mille dell'Arco predétto: & chi nel consumáre di quésto résto procedéssi con quéllo órdine,

Hugging Face — MODEL

metamorphoses-collection /
GiambullariOCR-1B-lora

2.48% character error rate

on held-out pages of a 1544 treatise never digitized

CER vs. off-the-shelf (lower is better)

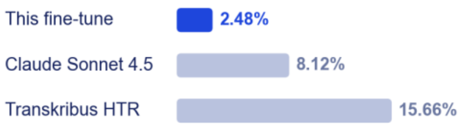

Base: LightOnOCR-2-1B · LoRA r32 / α64

16 training pages · ~15 min · 1× H100

Full 151-page book OCR'd in one day, solo

Model on Hugging Face · data + paper on Zenodo (CC-BY)

WHY RUNPOD, NOT A HYPERSCALER

RunPod
the AI developer cloud

- **Minutes to procure** — pay-per-use — this fine-tune cost $6.96
- **PyTorch + CUDA ready** — AI models run out of the box
- **Full SSH access** — an agent drives the box directly

TO THE LIMIT: LITERARY TEXT ANALYSIS ON A QPU

Pasqal

I wanted to test the claim as far from quantum computing as you can travel — the structural analysis of literary texts.

It doubles as a use case that is more accessible than the usual hardware demos — Z_2 phases in Quantum Magnets, or sampling for Machine Learning.

All of it on Pasqal Cloud credits — about 3,660 hours of cloud compute, 114 of them on a real quantum processor (from Sherbrooke, Canada), every job triggered from a Claude chat window.

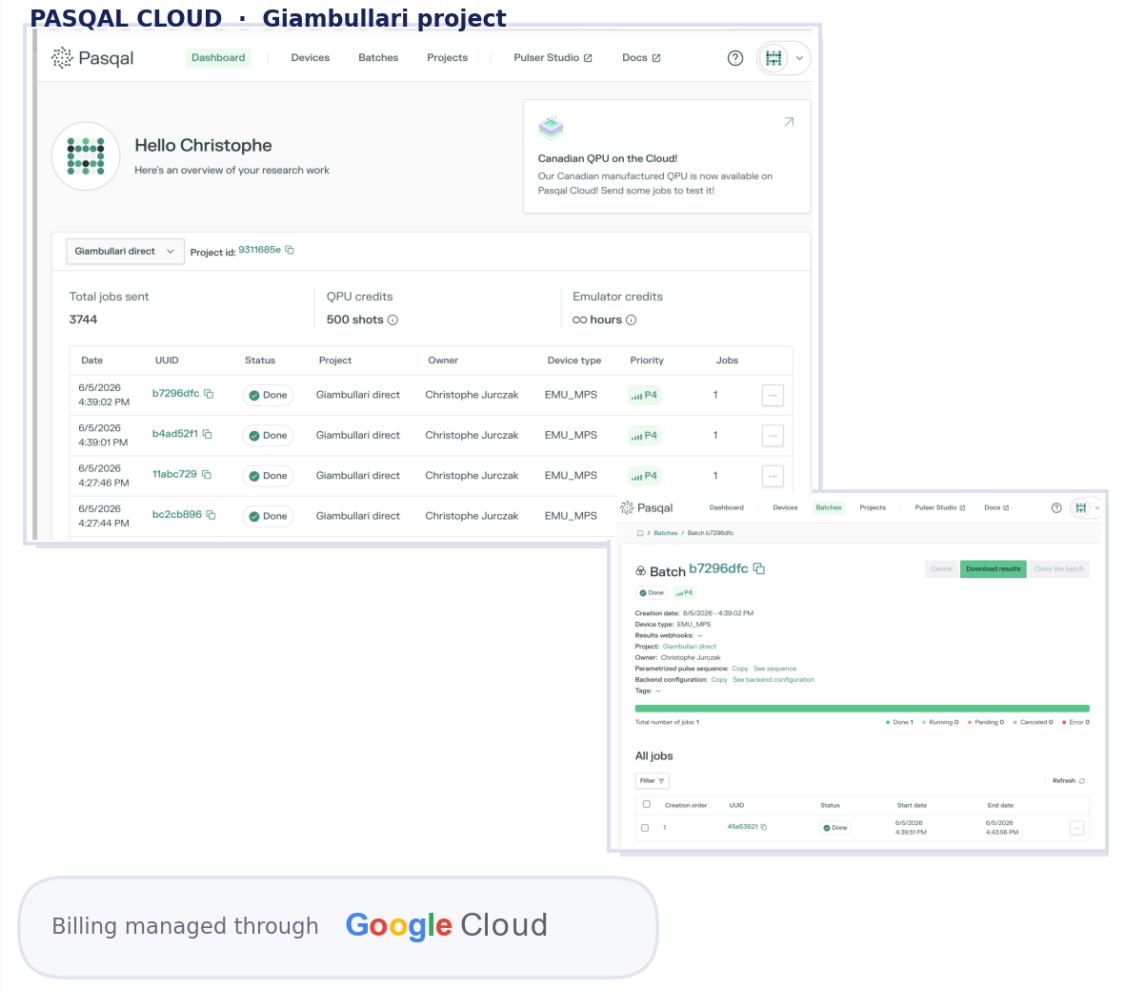

WHAT WE'RE ACTUALLY PROGRAMMING

1 · Place atoms = your graph

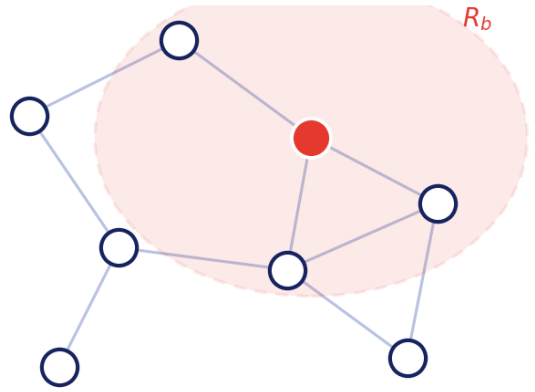

Atom positions encode the problem.
Reach within R_b = a constraint (an edge).

2 · One laser pulse = the program

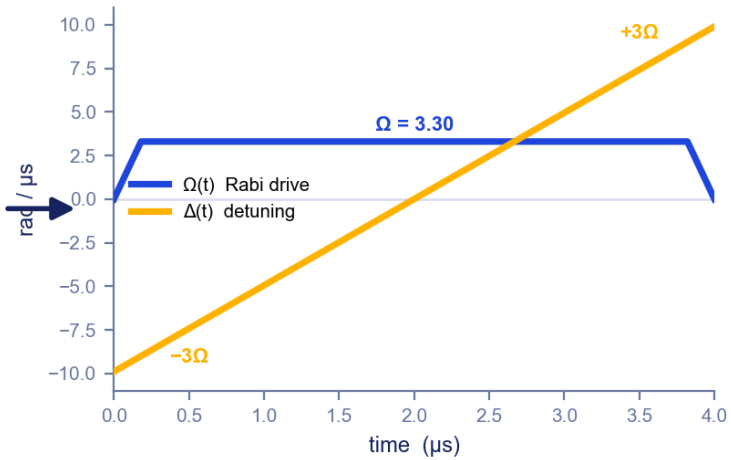

A single global sweep, ~4 µs.
No gates, no compilation.

3 · Physics relaxes to the answer

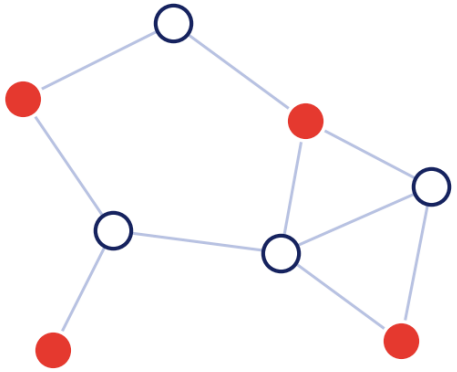

Lowest-energy 'on' set (red) =
Maximum Independent Set.

$$H(t) = \underbrace{\frac{\Omega}{2}\sum_i \sigma_i^x}_{\text{laser drives each atom}} \underbrace{- \Delta(t)\sum_i n_i}_{\text{detuning rewards 'on'}} \underbrace{+ \sum_{i<j} \frac{C_6}{r_{ij}^6} n_i n_j}_{\text{blockade penalises neighbours both on}}$$

the register IS the graph,
encoded in matter — not a simulation of it.

No gates, no circuits: you arrange atoms and sweep one laser, and iterate. The hard part is choosing what to compute and decoding the bitstrings that come back.

THE USE CASE: STRUCTURE READ AS PHYSICS

A book becomes a register of atoms. The *Maximum Independent Set* is its structural backbone; *rigidity* measures how unique that backbone is; you read it straight off the atoms.

Then I inverted it — wrote books to fit graphs the hardware likes. OuLiPo became QOuLiPo: 37 texts on Pasqal's FRESNEL, up to 100 atoms.

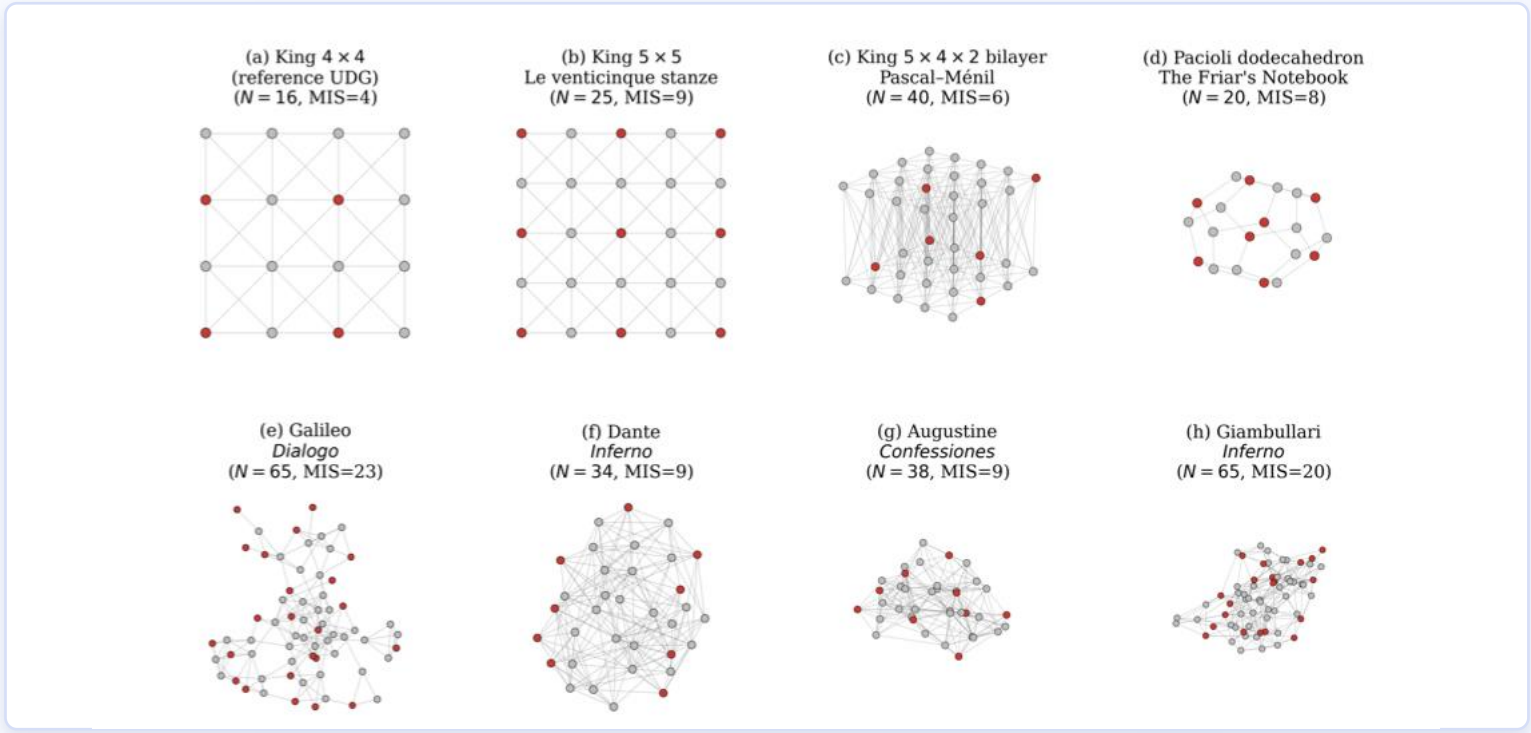

Engineered QOuLiPo graphs (top row) and natural-text graphs (bottom). Red nodes = the backbone the hardware finds.

REPRODUCE IT YOURSELF: FROM BOOK TO QPU

The agentic loop, five skills:

1 · paper-to-spec read the book's graph (N=25 · exact 2-D UDG)

2 · spec-to-sequence SA-embed to FRESNEL · build the pulse

3 · validate-emu EMU_MPS check — before any shots

4 · qpu-submit FRESNEL · 500 shots ≈ 3 min

5 · harvest histogram vs register-MIS · register plot

5×5 king register · 5 µm spacing · adiabatic −3Ω → +3Ω · 500 shots

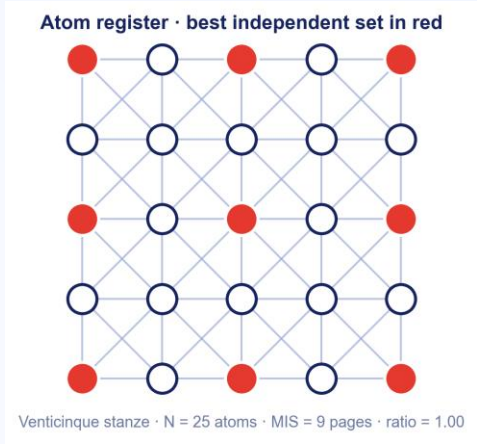

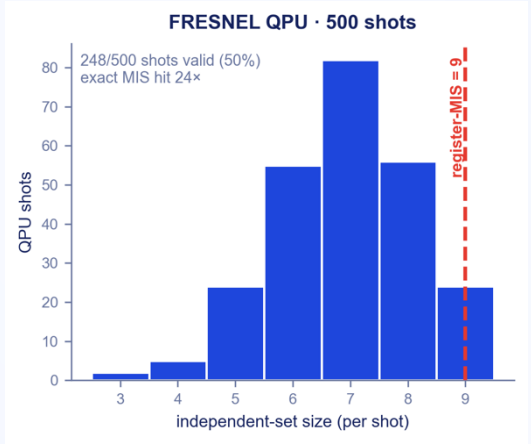

Anyone with a Pasqal Cloud account can rerun this — runbook and Pulser scaffold included.

LOWERING THE IMPLEMENTATION BARRIER — WITH PASQAL

1 · Read & extract arXiv PDF → Hamiltonian

2 · Build sequence Pulser · QUBO

3 · Validate on emulator EMU_MPS, noiseless

4 · Calibrate device noise FC1 noise model

5 · Run on QPU FRESNEL FC1

6 · Analyse QPU vs theory

Human approves at each gate.

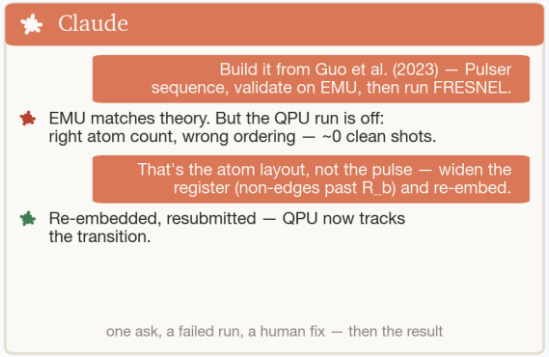

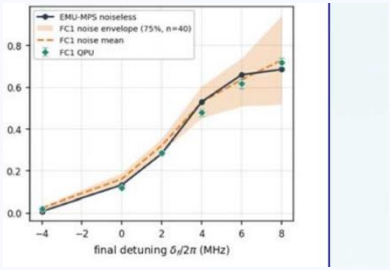

Reproduced Guo et al.'s (2023) triangular-lattice ordering transition on FRESNEL — QPU tracks theory.

49% of 526 Rydberg-array papers are runnable on today's hardware

Class	Papers	%
Implementable today	225	42.8
+ modest adaptation	33	6.3
Needs XY encoding	141	26.8
Needs local addressing	87	16.5
+ 4 smaller categories	*40*	*7.6*

The other half names the exact hardware to build next.

Pasqal

THE DEEPER BARRIER: TOO FEW ALGORITHMS TO RUN

Lowering the implementation barrier assumes you already know what to run. The harder problem sits upstream.

30+ years after Shor

still almost no quantum algorithm with a proven super-polynomial advantage and a confident use case outside cryptography. Hardware is racing ahead of the algorithms themselves.

Wide open

Only a handful of quantum algorithms have a proven advantage today. Discovering the rest is the field's biggest open problem.

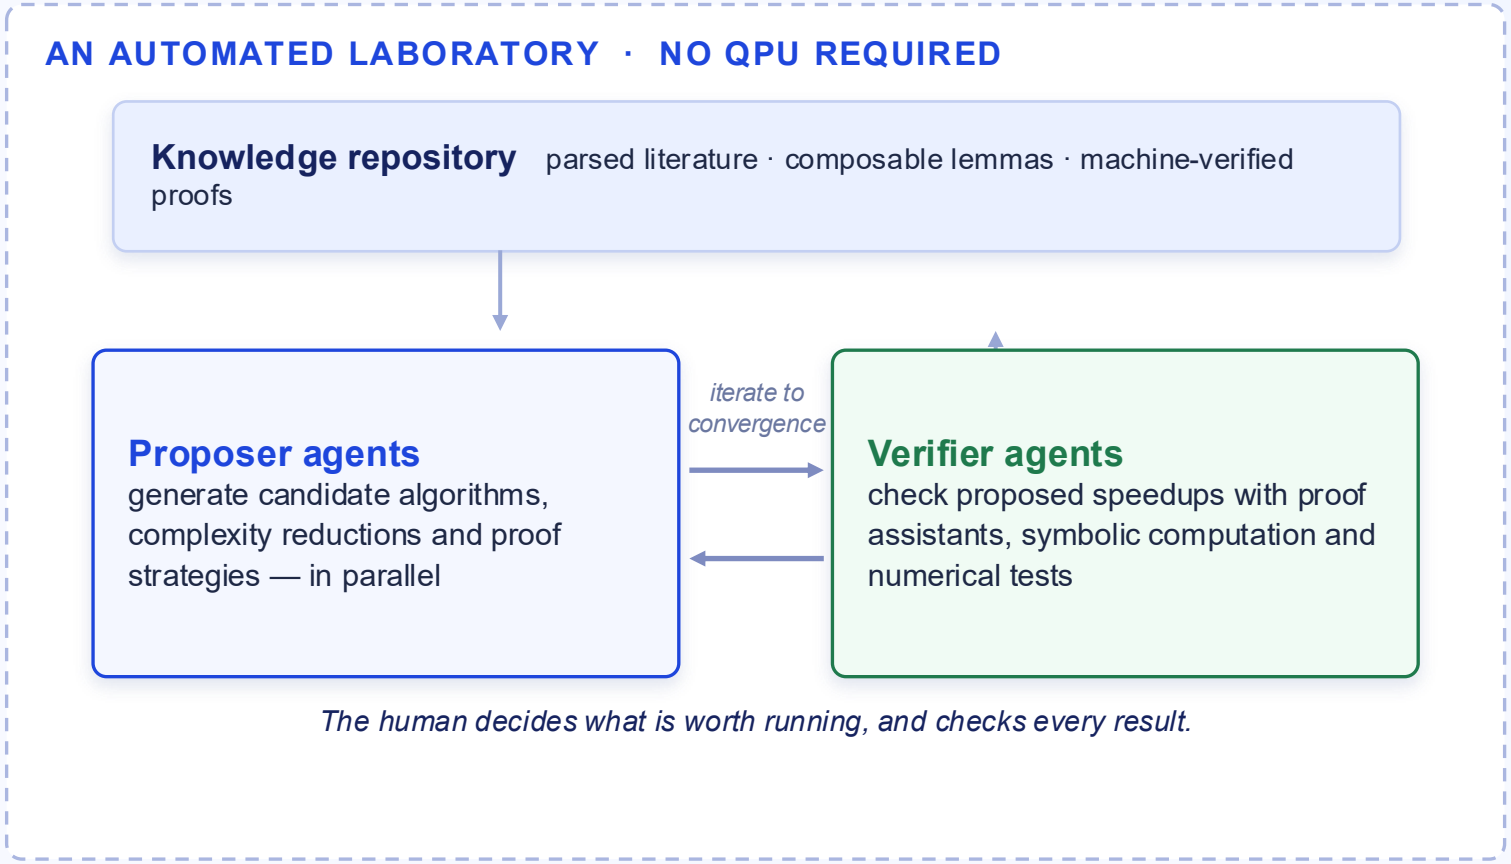

The bottleneck is verification: the infrastructure that makes AI output certifiable.

WHAT MATTERS

Invent the RunPod of QC

What we need from QPU vendors: reliable hardware · excellent noise models & calibration · documentation · libraries · primitives — built for machines, not papers. Access, not qubit count, is the gate.

Agents aren't geniuses

They clear the plumbing, not the judgment. Left alone, ours trusted the easy readout — the right number of atoms lit up, but not one was a valid answer. Which measurement matters is still on us. (NatureBench, 2026.)

Conversation, not a button

The push-button dream — hand over the task, get a coffee, come back to a solution — is wrong for hard problems. The value is the human–AI dialogue, with verification.

Each generation reads its world through the tools it has — ours just got a new instrument.