
Hybrid classical–quantum generative models for synthetic medical data

Soraya Panambalom

Quantum Group, School of Computing, and the Newcastle Solid Tumour Research Group

Dr Jonte Hance, Dr Dan Williamson, Dr Ed Schwalbe, Dr Nick Chancellor

Hardware access through the SparQ programme of the NQCC.



The cohort: 333 patients, Group 3 and 4 medulloblastoma

333

patients

16 • 5%

carry the MYC amplification


31

variables recorded per patient


103

deaths, where a survival model needs 300

The features that matter most clinically are
the ones the data can least support.



A tabular variational autoencoder with a quantum latent



17


models in all

5

things varied, one at a time

1

encoder, decoder and objective, unchanged throughout



The hardware runs: 4 configurations, 200 submissions

\$2,500

SparQ allocation on Amazon Braket

108


qubits, Rigetti Cepheus

200

submissions, 20,000 synthetic patients

3

runs of each configuration: exact, noise model, processor



Billing, shot budget, and verification

BILLING

\$0.11 not \$0.41

per patient, once we measured how Bracket charges

\$2,212 not \$8,125

for the same 20,000 patients, against a \$2,500 allocation

\$6.44

the two test jobs that settled it

SHOTS

250

shots, past which nothing moves by more than 0.4 of the noise

2,000

tested, and it bought nothing

CHECKS

8

checks, and 8 failures injected on purpose to test them

5 of 8

caught before a cent was spent

The model comparisons are in a paper under review.
Full protocol, the metric panel, and the results.
Preprint to follow.

Fitting the noise rate on the device's own output

0.53 – 0.99

readout fidelity across the chip, same operation qubit to qubit

0.50 – 0.99



two-qubit gate fidelity across the chip

9% – 43%

four answers from one formula on one machine

17 of 24

comparisons where the real processor beat the noise model



Postprocessing data loss in SamplerQNN

Also from the SparQ programme, on IBM processors, on a breast cancer task.

4 among 156

qubits placed by the compiler, with all 156 bits returned

It looks exactly like a barren plateau.

Near-constant loss, very small gradients, an optimiser that never learns. The cause is a filter in classical software, not the hardware.

up to 99%

of valid measurements silently discarded by one filter

**Reported as issue #1040, fixed in PR #1041,
merged May 2026.**

Written up with colleagues at Newcastle and at the NQCC,
including one of the Qiskit Machine Learning
developers: <https://arxiv.org/abs/2609.05060>

94% or 39%

accuracy from the same raw data, depending only on that step

Thank you.

Hardware access through the SparQ programme of the NQCC.

Cohort from the Newcastle Solid Tumour Research Group.

Soraya PANAMBALOM · Newcastle University