A Reproducible Evaluation Protocol for Quantum Non-Linearity in Variational Quantum Models

Which way of making a quantum circuit non-linear works best depends on what you measure

Pavel Sulimov Claude Lehmann

Zurich University of Applied Sciences (ZHAW), Winterthur · University of Zurich

NeurIPS 2026, Track on Evaluations and Datasets

A quantum circuit has no activation function

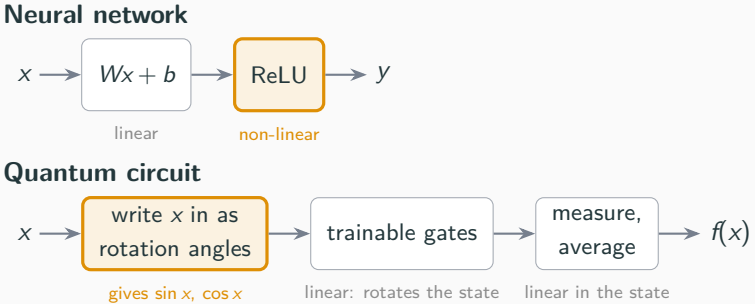

The output is a short Fourier series in x. Which frequencies it can use depends on how the input is fed in, so non-linearity has to be designed into the circuit.

Six published ways to design it in

	what it does	closest ML idea
Data reuploading	feeds x in again before every layer	every layer sees the input
Ancilla tracing	x enters once; a helper qubit is coupled in, then discarded	a deliberately lossy layer
Ancilla reuploading	both of the above	
Quantum CNN	convolution and pooling over qubits	a CNN on qubits
Strongly entangling	the standard dense circuit; x enters once	the baseline
Quantum reservoir	fixed random circuit; only the readout is trained	an echo-state network

How a new design is usually judged

one accuracy on one toy dataset, at a width, depth and training budget the authors chose

Four questions, asked at the same width, depth, budget and seeds

Frequency reach (NLC)

Can the trained model fit $\sin(kx)$ as k rises to 16?

ML cousin: spectral bias

Noise robustness (NSS)

How much of the trained fit survives gate noise up to 10%?

ML cousin: robustness to hardware noise at inference

Gradient health (1−BPPI)

How fast do gradients vanish as qubits go from 2 to 8?

ML cousin: vanishing gradients, here exponential in width

Learning efficiency (TER)

How fast, and with how few parameters, does it learn one-step forecasting of the chaotic Lorenz system?

ML cousin: sample and parameter efficiency

Each score lies in $[0, 1]$, higher is better, and ships with the raw curves it was computed from.

One composite score, with weights fixed before any result

$$\text{Composite} = 0.35\ \text{frequency reach} + 0.25\ \text{gradient health} + 0.20\ \text{learning efficiency} + 0.20\ \text{noise robustness}$$

(NLPS in the paper)

Why publish one

people will build a leaderboard anyway; a weighting frozen in advance means no author, including us, picks weights after seeing the numbers

How to read it

always next to its four parts; close values count as a tie; compare only circuits of the same width and depth

No circuit leads on every score

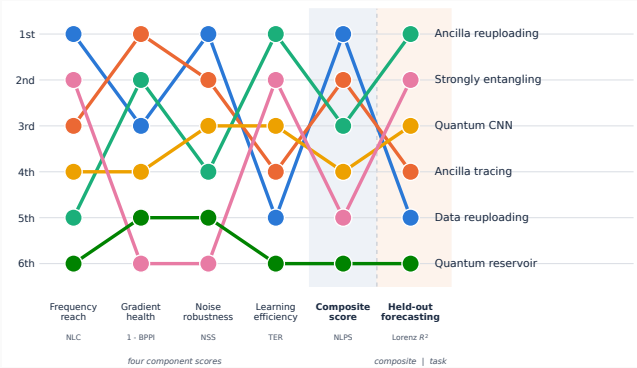

Data reuploading
feeds x in at every layer

Ancilla tracing
x once; helper qubit, discarded

Ancilla reuploading
both of the above

Quantum CNN
convolution + pooling on qubits

Strongly entangling
standard circuit, x once (baseline)

Quantum reservoir
fixed random circuit + readout

Rank of each circuit under each score at 3 qubits and 3 layers, mean of five seeds.

The composite winner is fifth of six on held-out forecasting

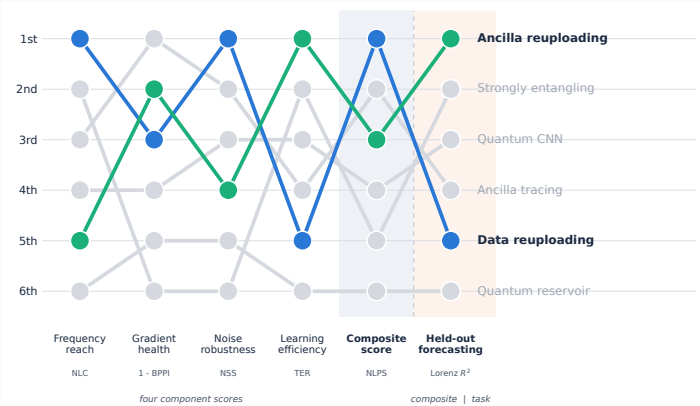

Data reuploading wins the composite through frequency reach, then forecasts worse than predicting the mean ($R^2 = -0.12$). **Ancilla reuploading** is third on the composite and forecasts best ($R^2 = 0.48$).

What the composite can and cannot tell you

composite vs. held-out task

0.03

rank correlation, essentially zero

test on your own task too

10^5 other weightings

0.97

agreement with the published ranking; top two never move

the weights are not the problem

3 $\rightarrow$ 4 qubits and layers

0.60

agreement between the two rankings; the leader changes

model size is: always match it

Spearman ρ (left) and Kendall τ (middle, right) over the six circuits, five seeds each.

Report the parts, at a matched model size

Take-away

Report the four scores next to the composite, at a stated and matched width and depth. Where the scores disagree, that disagreement is what a reader needs to see.

Try your own circuit

1. register one circuit-builder function
2. rerun the four stages; seeds and budgets are fixed
3. compare against a baseline at the same width and depth

About 2 hours on an 8-core laptop CPU, no GPU.



`github.com/sulimovp/quantum-nonlinearity-eval`

Poster 139813, Paris, 9 December 2026