VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER





QUANTUM RESERVOIR COMPUTING

FOR FRAUD DETECTION

Ivana Miháliková, Renáta Rusnáková, Van Binh Henri VU

Fraud can be hard to spot in a single transaction

Every individual payment can be perfectly ordinary. The signal lives in **the pattern** across weeks of activity.

MAIN CHALLENGES IN DETECTING MONEY LAUNDERING:

Scale

Investigators cannot review every transaction manually.
The model helps identify accounts that need attention.

Rarity

Fraud cases are uncommon.
Too many false alerts would overwhelm investigators.

Time

The order and timing of transactions may carry useful information.
We need to examine the sequence, not just isolated payments.

So the unit of analysis is the account over time, not the individual payment.



From transaction history to an account-level decision

INPUT:
One account, a variable number of transactions



a handful of transactions



hundreds of transactions



thousands of transactions

OUTPUT:
One prediction for the account

fraudulent

legitimate

Sequence lengths **differ** by orders of magnitude, but a classifier needs a fixed-size input.
Something has to turn a whole history into **one vector**.



Sequence models: what they offer and what they require

Train a recurrent model

- LSTM, GRU or a Transformer over the sequence
- Learns relationships between events over time
- State of the art when data are plentiful

What you pay for it

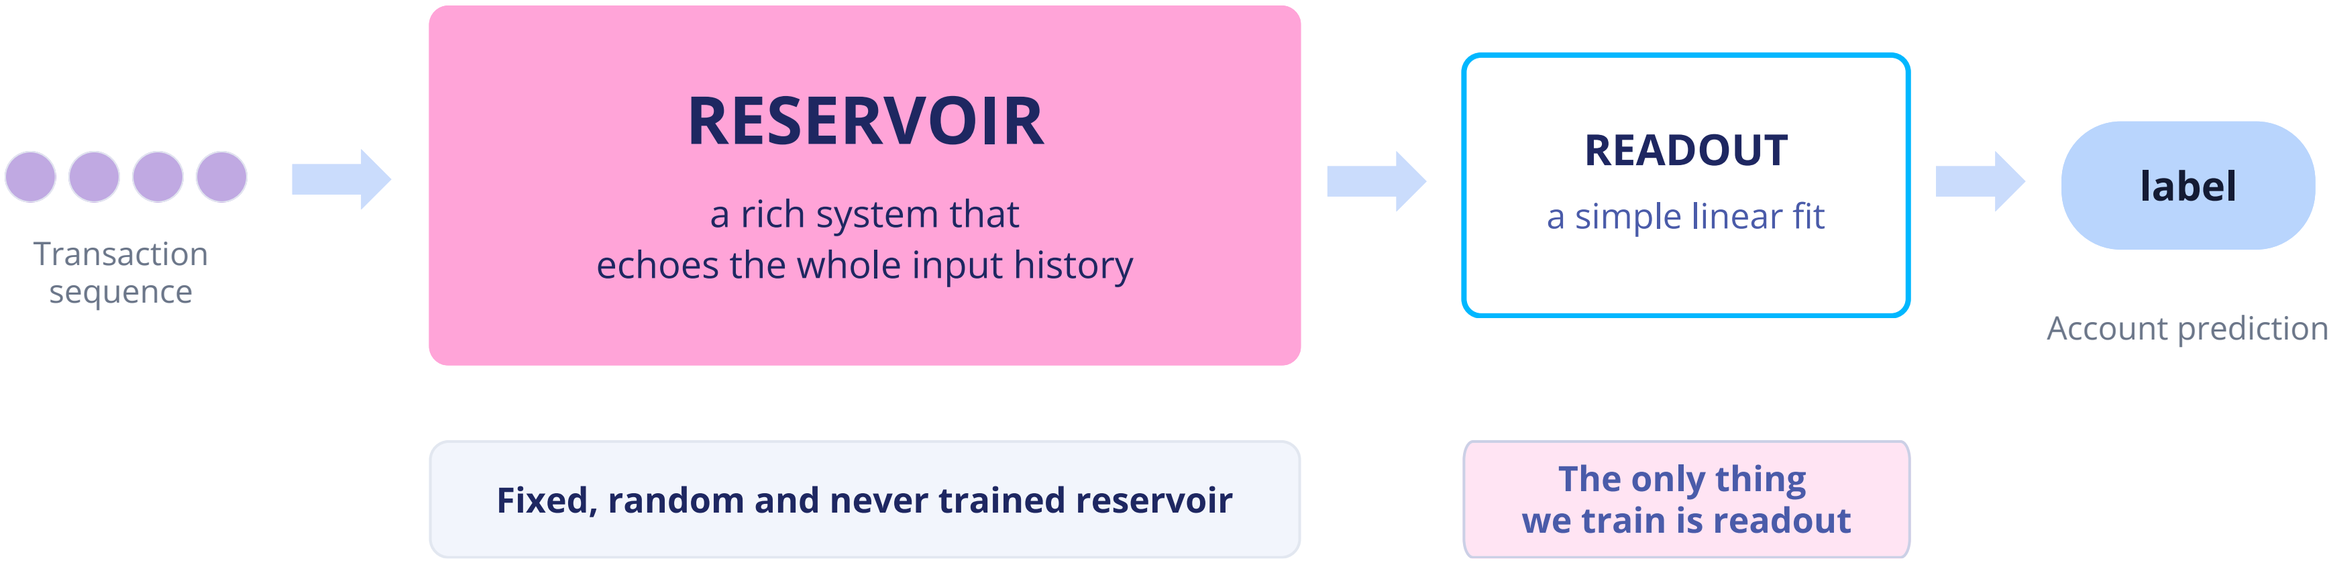
- Thousands of parameters trained through time
- Expensive to retrain as behaviour drifts

What if we could process the sequences without training a large and complicated model?



Reservoir computing

In neural networks, we train many paramters to learn patterns in data.
Here, we feed the sequence through a fixed, random system and train only a linear readout on what comes out.



The reservoir must retain enough information to distinguish different transaction histories.

Only the final linear readout is trained, making the training process simpler and more efficient.



Why use a fixed reservoir?

Training

The reservoir and quantum circuit remain fixed. We train only the final linear readout, so there is no backpropagation through the transaction sequence and no circuit gradients to estimate.

Limited fraud labels

Confirmed fraud cases are difficult to obtain and may be labelled only after an investigation. Training fewer parameters is useful in this setting, although it does not guarantee good results with a small dataset.

Controlled comparisons

We can run the quantum model and its classical control with the same data split and random seed. This helps us isolate the effect of adding the quantum measurements.

The question then becomes: what does the quantum circuit add?



What does the quantum circuit add?

8 QUBITS
=
A 256-dimensional state space

Each additional qubit doubles the number of state dimensions.

Features from measurements

Transaction values set the angles of rotation gates. After applying the circuit, we measure individual qubits and correlations between pairs of qubits. These measured values become features for the classical reservoir state.

Nonlinear dependence on the input

The circuit acts linearly on a quantum state. However, we encode input features into rotation angles, and the measurement results depend nonlinearly on these angles. This allows the quantum circuit to generate nonlinear features from the input data.

We keep the circuit fixed and test whether its measured features improve the account-level prediction.



Hybrid quantum reservoir

One pass per transaction. The circuit supplies features, a classical state carries the memory.

1

Encode

Map the transaction features to rotation angles.

2

Evolve

Apply rotation and entangling gates to the qubits.

3

Measure

Measure each qubit and pairwise correlations.

4

Integrate

Update the classical reservoir state using the new measurement features.

5

Summarise

The whole series collapses to one fixed-size vector. The readout predicts the account label.

The circuit and reservoir weights stay fixed. We train only the linear readout using account labels.



How do we test the quantum contribution?

- We compare the hybrid reservoir with a matched classical control.
- Both use the same data splits, random seeds and readout.
- The hybrid model also receives features from quantum measurements.

Quantum reservoir

classical state + quantum measurements

The circuit contributes extra features at every step of the sequence.

Matched classical control

classical state only

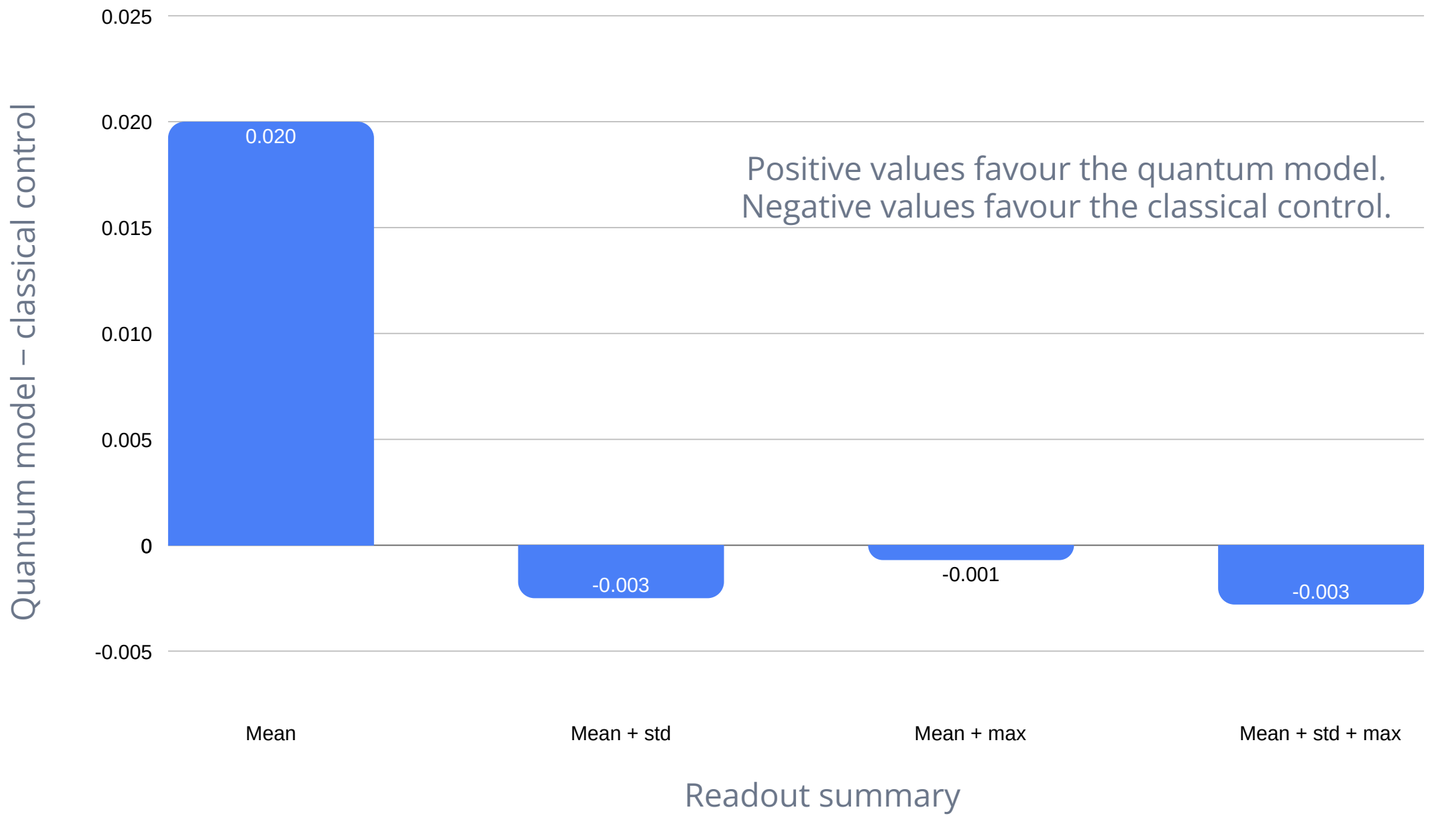
Same random matrices, same seeds, same splits, same readout, same metric. One term removed.

- Each pair is evaluated on the same split and seed.
- No account appears in both training and evaluation data.
- We run 100 paired comparisons per configuration.



Preliminary results

The result depends on the readout



SIMPLE READOUT

The quantum model performed better in 89 of 100 paired runs.

RICHER READOUTS

The quantum model performed better in 32–37 of 100 paired runs. The classical control performed better more often.

PR-AUC (Area Under the Precision-Recall Curve) - this metric is useful when one class is much less common than the other, as is often the case in money laundering detection. Precision tells us how many of the accounts flagged by our model are actually positive cases.

The quantum features improve performance with simple pooling, but this advantage disappears when richer classical features are used.



Summary

- Reservoir computing offers an interesting way to analyze sequences without training a large and complicated model.
- By adding quantum measurements, we can introduce additional features into a classical reservoir.
- Our preliminary results show that these quantum features can improve performance in specific configurations.







THANK YOU

FOR YOUR ATTENTION!