Objectives for this session:

To traverse the amazing high-dimensional spaces of quantum models

What is Hilbert space
What is quantum evolution
Optimisation of quantum models in quantum and classical spaces
The curse of dimensionality

- *Naive n-ball metaphor*
- *Barren plateaus (too many parameters)*
- *Orthonormal desert (too many qubits)*
- *The statistical paradox (too many measurements)*

From naive n-balls to quantum manifolds
Hilbert space is not a n-ball !
Getting the intuition right!
Summary

Paradoxes of High-Dimensional Quantum Models

and... why is it so damn hard to train them !?!

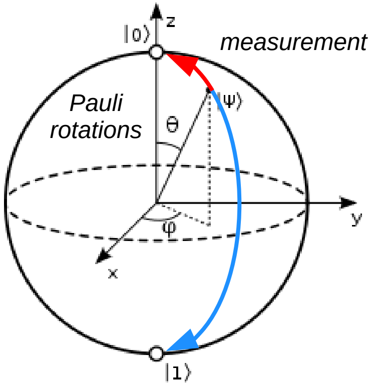

Jacob L. Cybulski
Enquanted, Australia

June 27, 2026, Fundacja Quantum AI, QML Conference 2026.

Nielsen, M., Chuang, I., 2010. *Quantum Computation and Quantum Information: 10th Anniversary Ed*. Cambridge University, New York. https://doi.org/10.1017/CBO9780511976667.

The Hilbert Space
and the quantum state evolution

```python
dev = qml.device('default.qubit', wires=2)

@qml.qnode(dev)
def project_to_3d_v1(params):
    qml.RY(params[0], wires=0)
    qml.RY(params[1], wires=1)
    qml.CNOT(wires=[0, 1])
    qml.RZ(0.5 * params[0] * params[1], wires=0)
    qml.RZ(0.3 * params[0], wires=1)
    return (
        qml.expval(qml.PauliX(0)),
        qml.expval(qml.PauliZ(1)),
        qml.expval(qml.PauliZ(0) @ qml.PauliZ(1))
    )
```

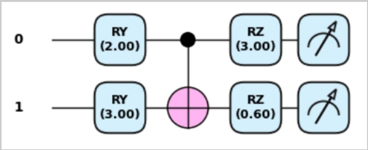

Consider the above quantum model.

As a quantum model executes, its *state evolution* is given by a *path*, or a "trail" left behind by *the state vector*.

Because of entanglement, the state vector is a "complex" arrow and the trail becomes "knotted".

Because of the Heisenberg uncertainty and noise, each time the model executes, it takes a slightly different path.

Trainable parameters allow adjusting the direction we take.

By varying encoded data and model parameters, all possible state vector paths create a smooth curved surface – a *manifold*.

The Hilbert space is a coordinate system and a space of directions that the state of your model could take during its execution.

Notation:

$p_0 \equiv \texttt{params}[0],\ p_1 \equiv \texttt{params}[1],\ c_i \equiv \cos(p_i/2),\ s_i \equiv \sin(p_i/2),\ \phi_0 \equiv p_0 p_1/4,\ \phi_1 \equiv 0.15\, p_0.$

Gate-by-Gate state evolution and observable projections

Step	State / Observables	Comment
1	$\lvert\psi_1\rangle = c_0\lvert 00\rangle + s_0\lvert 10\rangle$	$R_Y(p_0)$ rotates qubit 0 into superposition.
2	$\lvert\psi_2\rangle = c_0c_1\lvert 00\rangle + c_0s_1\lvert 01\rangle + s_0c_1\lvert 10\rangle + s_0s_1\lvert 11\rangle$	$R_Y(p_1)$ rotates qubit 1; state remains separable.
3	$\lvert\psi_3\rangle = c_0c_1\lvert 00\rangle + c_0s_1\lvert 01\rangle + s_0s_1\lvert 10\rangle + s_0c_1\lvert 11\rangle$	CNOT swaps $\lvert 10\rangle \leftrightarrow \lvert 11\rangle$; generates entanglement.
4	$\lvert\psi_4\rangle = e^{-i\phi_0}(c_0c_1\lvert 00\rangle + c_0s_1\lvert 01\rangle) + e^{+i\phi_0}(s_0s_1\lvert 10\rangle + s_0c_1\lvert 11\rangle)$	$R_Z(\phi_0/2)$ on qubit 0 applies control-dependent phase, adds nonlinear/interactive warping.
5	$\lvert\psi_{\text{fin}}\rangle = c_0c_1e^{-i(\phi_0+\phi_1)}\lvert 00\rangle + c_0s_1e^{-i(\phi_0-\phi_1)}\lvert 01\rangle + s_0s_1e^{+i(\phi_0-\phi_1)}\lvert 10\rangle + s_0c_1e^{+i(\phi_0+\phi_1)}\lvert 11\rangle$	$R_Z(\phi_1)$ on qubit 1 adds independent phase, linear/gauge-like warping; final state.
6	$\langle Z_1\rangle = \cos p_0 \cos p_1,\ \langle Z_0Z_1\rangle = \cos p_1$ $\langle X_0\rangle = \sin p_0 \sin p_1 \cos(p_0p_1/2)$	Expectation values: Z-terms phase-invariant; $\langle X_0\rangle$ encodes nonlinear interference, it contains $\cos(p_0p_1/2)$ because the R_Z rotation "drags" the state around the Z-axis, modulating the projection onto the X-axis.

Dirac "bra-ket" notation

Optimising a
PennyLane model

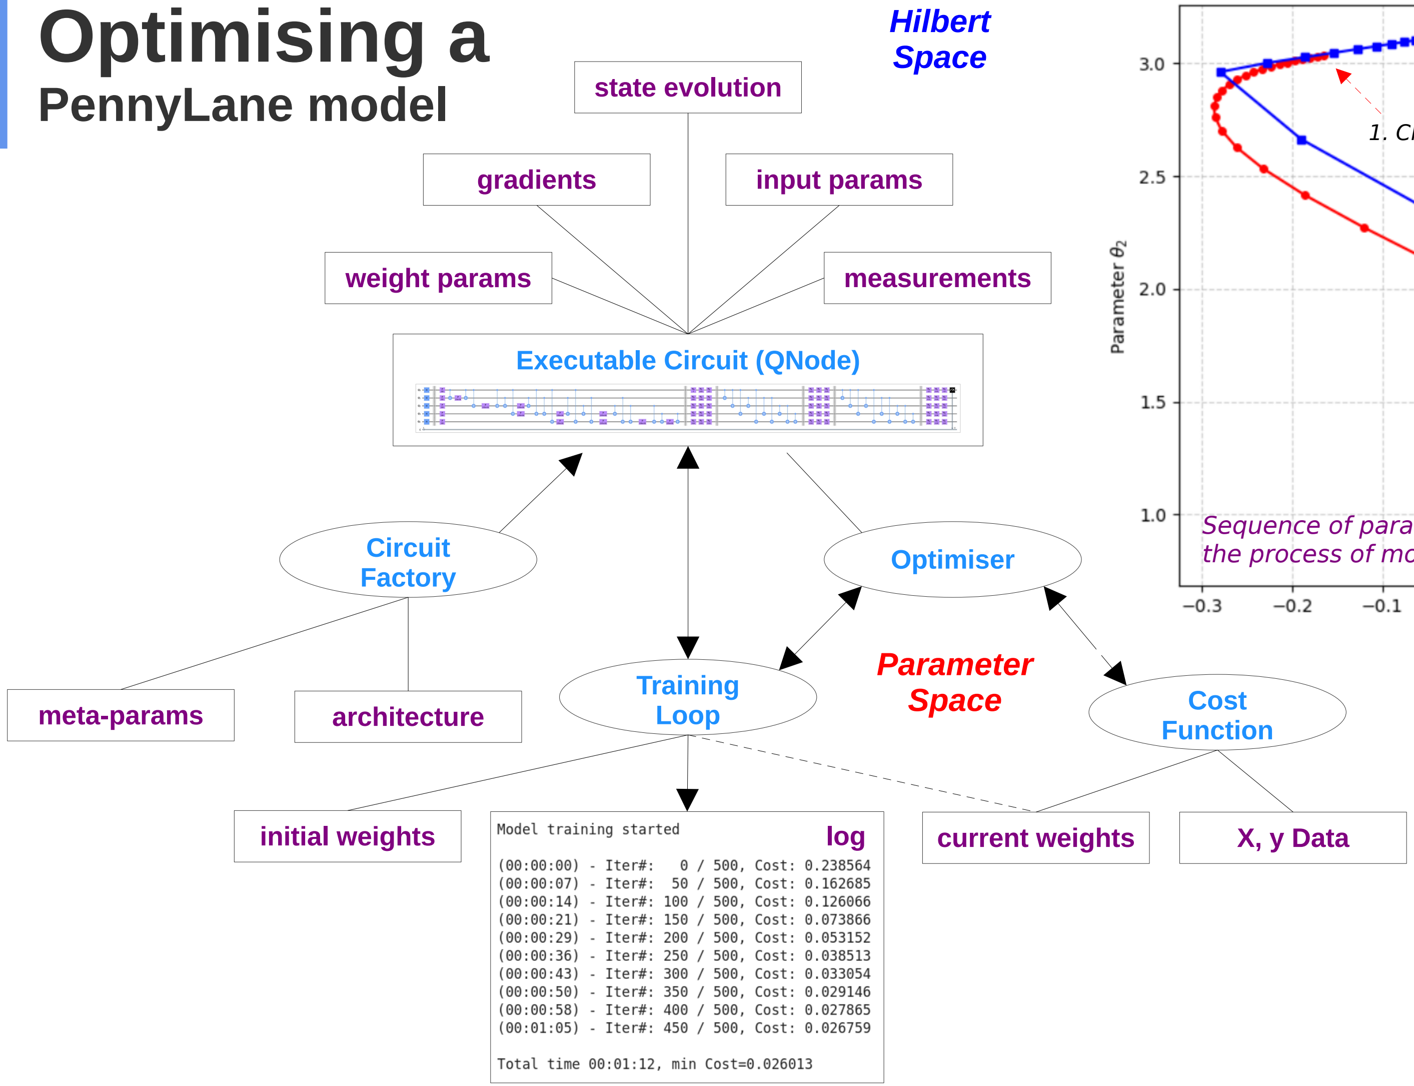

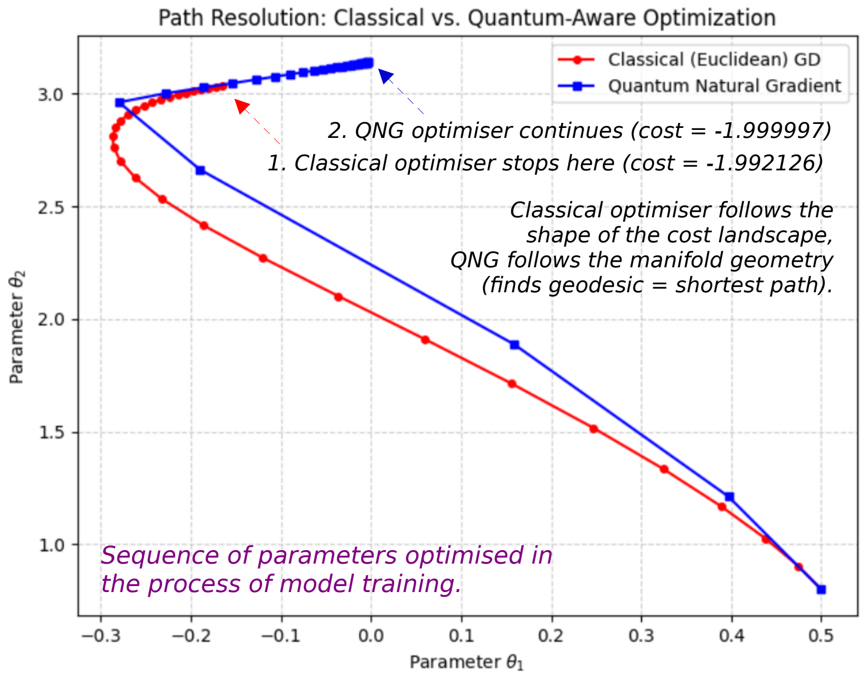

Hybrid quantum-classical VQA process

Parameter space

(dim = the number of params) is a classical multi-dim space of trainable gate parameters, which the optimiser navigates

Classical parameters

A classical optimiser can only manipulate the gate parameters but is not aware of their action on qubits states or their correlations via qubit entanglements.

Measurement

of individual qubits collapses their states, consequently projecting the circuit state onto classical outcomes - the only information reaching the optimiser. In the process we lose much of quantum information (e.g. phase).

Optimisation trajectories in quantum and classical spaces

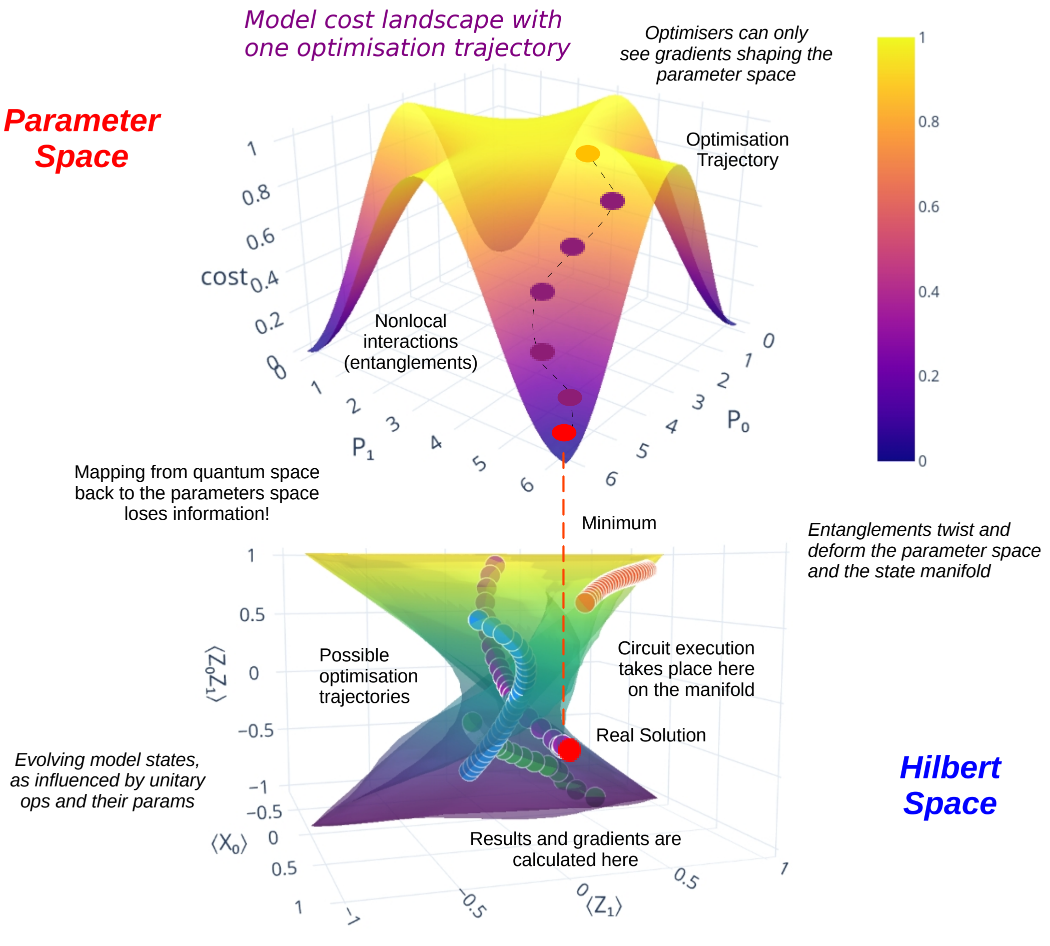

Warped model manifold with four classical optimisation trajectories (projection into the observables)

Hilbert state space

(dim ≈ $2^{\text{the number of qubits}}$) is the quantum realm where the models and their states evolve in response to unitary operations as defined by the circuit gates

Data encoding

brings in classical data into the Hilbert space as unique and correlated quantum states during the model execution

Entanglements

(CNOTs) create non-separable qubit states, which cause non-local interactions, and which warp the state space geometry, and as a side-effect, also the cost landscape

Circuit layers of gates

determine the evolution of the model's initial state into its final state during its execution (forming a state path)

The curse
of dimensionality

Jacob L. Cybulski, 2026. "A dance of the blind puppeteer: the interplay between a classical optimizer and the Hilbert space." in Oswaldo Zapata (Ed), *A Portrait of Quantum Technologies in Finance*, The Quantum Finance Boardroom, pp 107-121 (also preprint).

Cybulski, J.L., Nguyen, T., 2023. “Impact of barren plateaus countermeasures on the quantum neural network capacity to learn”, *Quantum Information Processing* 22, 442 (also online).

Barren Plateaus **(too many parameters)**

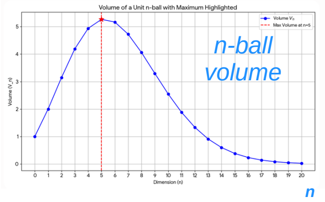

- Let us start with a naive, purely geometric conceptualisation of quantum phenomena.
- Consider a *n*-dimensional ball ≡ *n-ball* of *pure states only*, where all quantum phenomena are simplified.
- *n-ball volume shrinks with n* (for $R=1$, max @ $n \approx 5.25$, see above), derived from the volume of *n*-ball with radius $R=1$, which is:

$$V_n(R) = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)} R^n, \quad and \quad V_n(1) = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)}$$

and the *gamma function* $\Gamma(z)$ is defined as:

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t}\, dt, \quad \text{for } \mathrm{Re}(z) > 0$$

- When volume of the *n-ball of parameter space* is diminishing, the surface of the enclosed cost landscape flattens, consequently *all gradients are approaching zero* ≡ *barren plateau (BP)*.
- When BPs emerge, the optimiser struggles finding the minimum.
- *Random initialisation of weights makes it worse*, as random points may be far from the optimum, finding the minimum cost difficult!

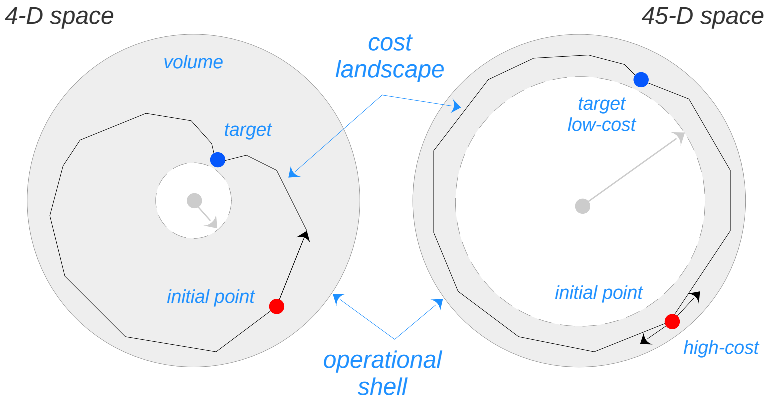

Yet, *the volume of a n-ball surface shell expands* (as a % of the total), as can be shown:

$$V_{shell}(\epsilon) = V_{total}\left[1 - (1-\epsilon)^n\right]$$

Where the shell thickness is $\epsilon > [0, 1)$ and as $n \to \infty$, the term $(1 - \epsilon)^n$ vanishes, for example:

- In 3-D (n=3), V_{shell} *(1%) ≈ 2.9%* of the V_{total}
- In 45-D (n=45), V_{shell} *(1%) ≈ 36.4%* of the V_{total}
- In 500-D (n=500), V_{shell} *(1%) ≈ 99.3%* of the V_{total} !

This is happening at just 500 parameters or 9 qubits (2^9 dimensions)!
Note that the parameter space is “suspended” over the Hilbert space!

The curse
of dimensionality

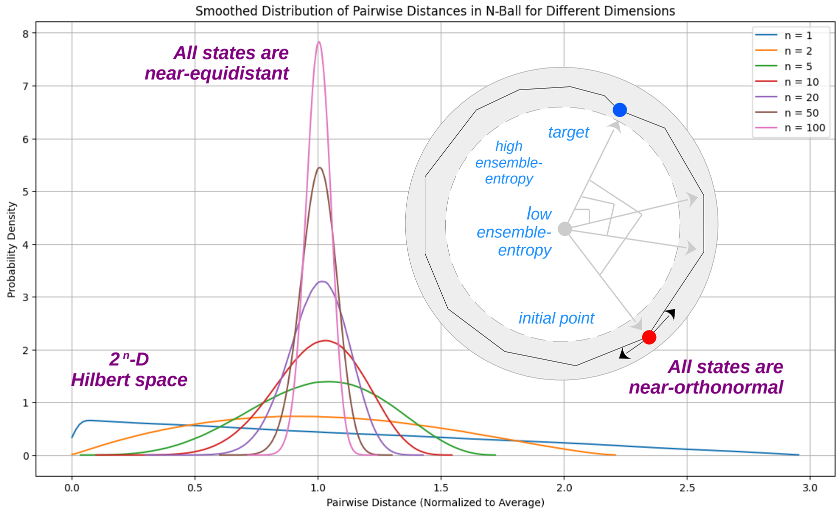

The concentration of angular separation around orthogonality in a n-ball can be formally justified by:

- in each coordinate dimension, the product of randomly selected state vectors is equally likely to be positive or negative;
- with raising dimensionality, the sum of these products converges statistically toward zero, leading to near-orthogonality between randomly selected state vectors.

Geometrical State Space Metaphor vs.
Orthonormal Desert (too many qubits)

- In the *n-ball of pure states*, distances concentrate around the mean ≡ *state distances are near-identical in all directions.*
- Any two random vectors pointing to the surface of the n-ball are near-orthogonal, creating the *orthonormal desert* of model states.
- In such a regime, it can be shown that statistically *all states are near-equidistant and orthogonal to the target*.
- For *n*-ball of $R=1$, the distance between a pair of states is:

$$dist(p1, p2) \approx \sqrt{2}$$

- During optimisation, the classical optimiser is is trapped on a high-dimensional manifold where the concept of *volume*, *distance* and *direction* has been erased by its geometry.
- *The path toward the target (solution) ceases to exist*.
- Also, the n-ball of pure states has a *strange entropy distribution*. Radial probability of finding a state at a given distance from the centre is $P(r) = nr^{n-1}$, the *n-ball centre is of low ensemble-entropy* and the *outer shell is a high ensemble-entropy* of states.
- Paradoxically, this leads to concentration of ensemble-entropy, where *uncertainty concentrates in a very narrow range of space*.

Here we are referring to coherent or ensemble-entropy of pure states based on distribution of their variants in the manifold, rather than thermodynamic entropy, where entropy of pure states is zero.

The curse
of dimensionality

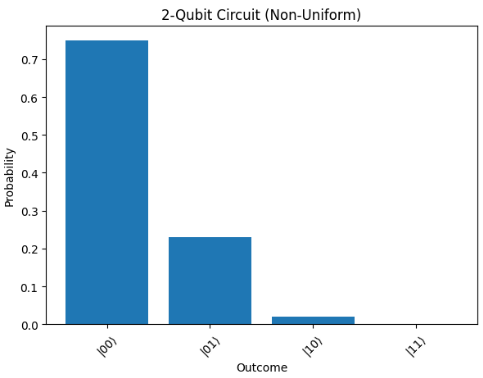

A) Few measurements = great precision and clarity

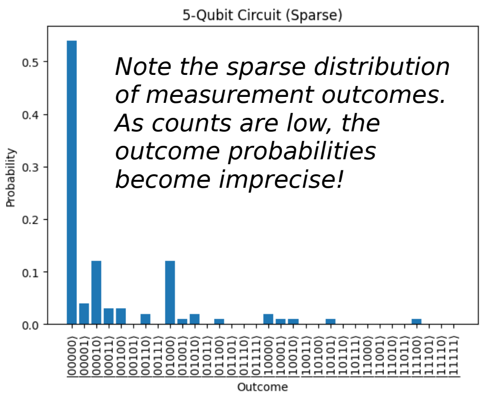

B) Many measurements = poor precision and clarity

Geometrical State Space Metaphor vs.
Sparsity of Measurements (too many measurements)

- When executing a circuit we need to repeat its measurements several times (the shots)
- It is common to measure all qubits (global measurements) as it allows extraction of complex correlations between qubits and provides rich feedback to the optimiser on the possible weight changes.
- However, when increasing the number of measurements, distribution of outcomes becomes sparse and the probability calculations become imprecise!
- *To ensure the same level of precision, we need to increase the number of circuit runs (shots) exponentially!*
- Alternatively, measure only a selection of qubits (local measurement).

All paradoxes presented here have been shown to exist in a naive, purely geometrical metaphor of the high-dimensional quantum state space.

However, these paradoxes are applicable to all high-dimensional objects and spaces, including the true representation of the projective Hilbert space.

From *naive* n-balls to *quantum* manifolds

Ingemar Bengtsson and Karol Życzkowski (2006):
Geometry of Quantum States: An Introduction to Quantum Entanglement, Cambridge Univ. Press.

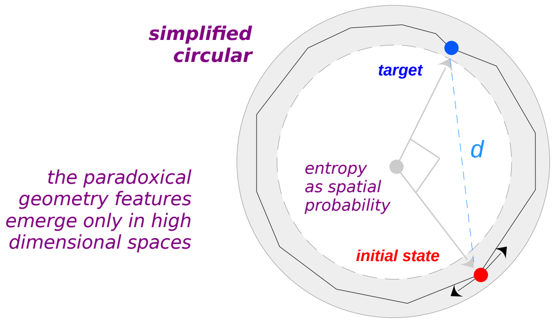

A) Visualisation of pure states in high-dimensional shell of naive n-ball space

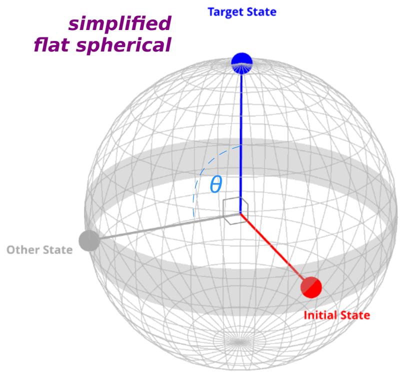

B) Visualisation of pure states in high-dimensional manifold of the Hilbert space

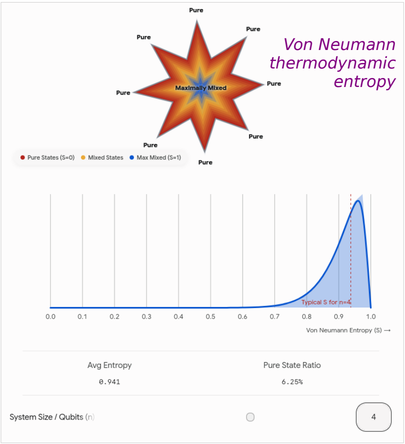

- **In the high-dimensional naive n-ball of pure states:**
  - all states are concentrated in a narrow shell near its surface (fig. A);
  - the space is flat and distance measures between states are Euclidean;
  - both the initial and target state are in this shell;
  - the vectors between pairs of states are near-orthogonal;
  - entropy represents spatial probability, high entropy near the surface.

- **In the high-dimensional quantum space:**
  - all states are represented within a complex multi-dimensional projective manifold in Hilbert space, shaped like asymmetric, spiked mace (fig. B), dictated by the constraint that probability of alternative states = 1.
  - the manifold is measured using Fubini-Study metric, i.e. as a quantum fidelity (overlap) between two states (derived from polar coordinates);
  - the interior of the manifold are full rank mixed states (all eigenvalues $\lambda > 0$), its outer boundary are states losing full rank (at least one $\lambda = 0$), the manifold of pure states is within the outer boundary (all but one $\lambda = 0$), pure states represent the extreme points of the outer boundary;
  - in high-dimensions, the angular distance between pairs of states is near-orthogonal, therefore from the perspective of any target state, all other states are concentrated in the massive equatorial area (Levy's Lemma);
  - the equator includes the initial and the majority of intermediate states;
  - low entropy states exist near the outer boundary (opposite to the n-ball);
  - all n-ball paradoxes manifest in this high-dimensional space as well!

What it looks like ...

not quite a n-ball

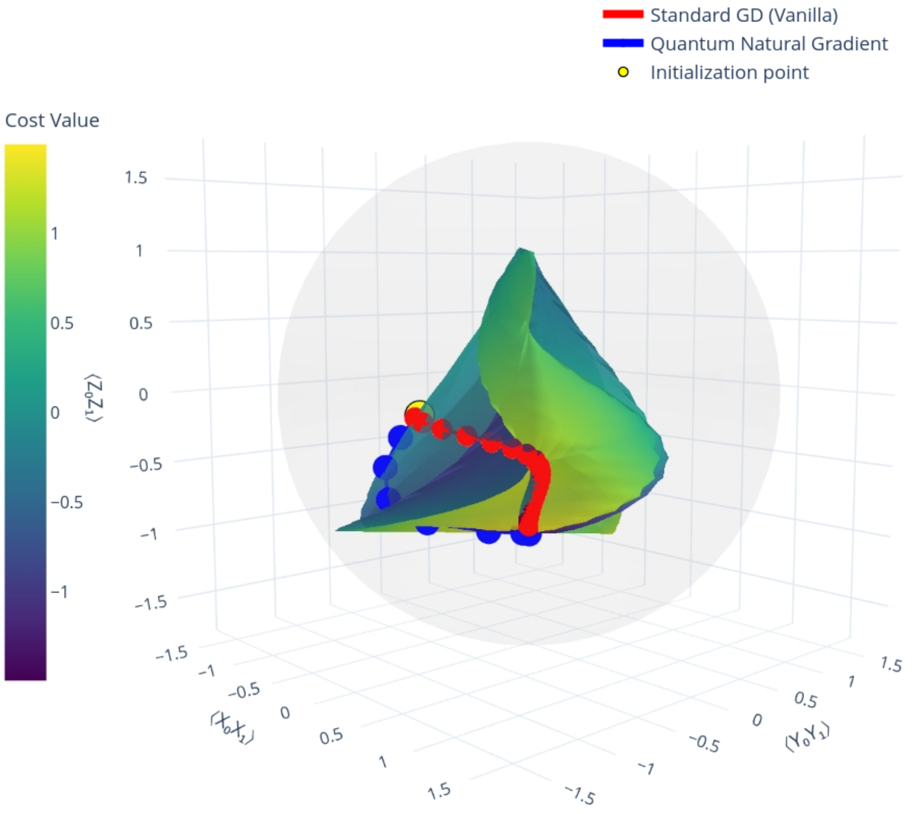

The selected cost function aims to simulate a VQE-like landscape.

Hilbert space is not a n-ball ...

- The canonical representation of a quantum model's manifold in the Hilbert space uses:
  - *Generalised Bloch body is its true representation*, where a *n* qubit system resides in space governed by the $SU(2^n)$ Lie group.
  - The state of the 2 qubit manifold decomposes into 15 non-trivial Pauli tensor products (e.g. $X \otimes I$, $Z \otimes Z$, etc.), so the generalised Bloch vector is 15 dimensional.

$$\rho = \frac{1}{4}\left(I \otimes I + \sum_{i=1}^{15} b_i \sigma_i\right)$$

  - The radius of the pure state manifold is linked to its purity $\mathrm{Tr}(\rho^2) = 1$, where:

$$\mathrm{Tr}(\rho^2) = \frac{1}{4}(1 + |b|^2)$$

  - We can project the model manifold into the Pauli axes, which are based on two-body correlation observables, so the plotting axes are (X_0X_1), (Y_0Y_1), (Z_0Z_1), rather than using mixed observables.
- When we create such a visualisation we can see that the geometry of the manifold is violently crumpled and the area of the lowest cost is incredibly warped as well (occluded dark blue).
- The classical gradient descent is unaware of the manifold geometry, so it traverses it in highly inefficient way.
- On the other hand, the Quantum Natural Gradient follows the shortest geodesic path along the manifold surface towards the lowest cost.

Getting the intuition right!

What can we see here?

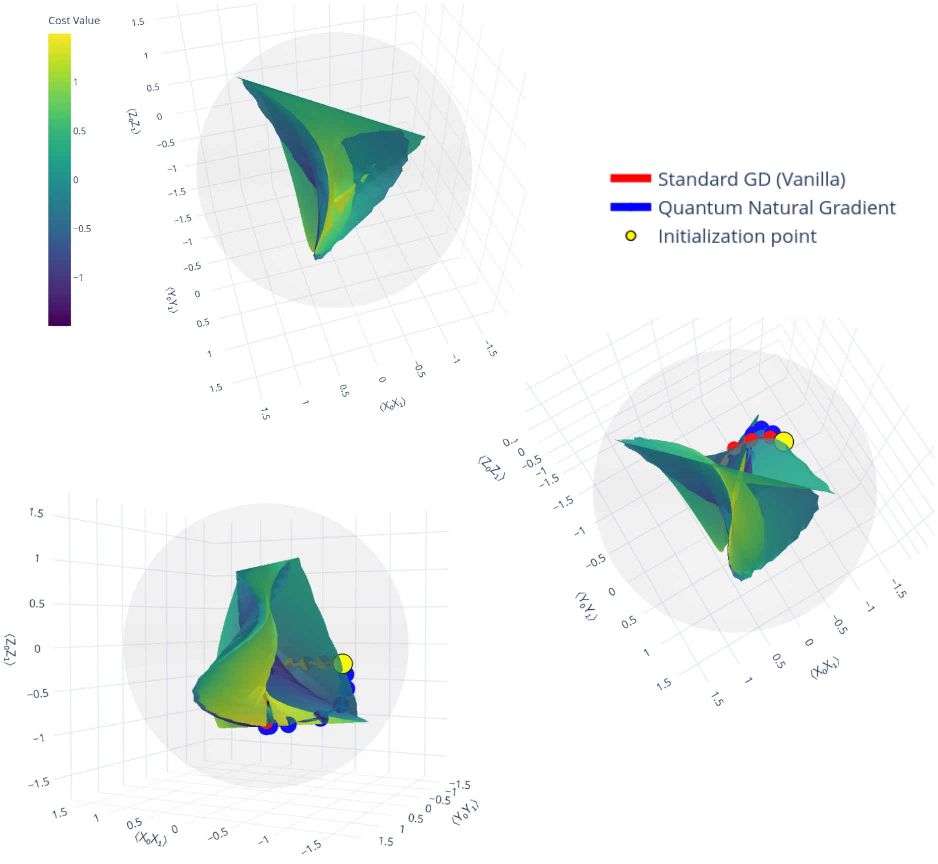

The selected cost function aims to simulate a VQE-like landscape.

Multi-dimensional multi-qubit manifold explained

- This visualisation identifies several features of the model's high-dimensional geometry:
  - The projection from 15 Pauli tensor products to the correlation subspace creates a geometry with the boundary of physically valid states forming a convex polytope **(here tetrahedron)**.
  - **The Information Trade-Off:** The 3D polytope maps the physical boundary between classical independence and global quantum correlation.
  - **Center Origin (0,0,0):** Maximum local entropy; where unentangled single-qubit states collapse.
  - **Faces & Edges (Surfaces):** Highly entangled, lower-rank mixed states constrained by non-linear physical boundaries.
  - **Vertices (4 Points):** Pure global information; the four maximally entangled Bell states.
  - Observed curvatures (including internal folds) are the result of the quadratic functions of trigonometric terms curving nonlinear trajectories in the observable space, and some are the artefact of the projection.
- While classical optimizers blindly scrape across the rank-3 faces to minimize loss, the Quantum Natural Gradient (QNG) tracks the true quantum information geometry, riding the rank 2 edges as natural geodesics (normally) toward the maximally entangled vertices.
- However, because of non-linearities in rotational operators and VQE-like cost function, the optimum point was found in the middle of the edge, which provided the ground state. In this case, vertices are outside the expressible capability of the ansatz.

Summary and Q&A

- Multi-dimensional geometry is full of paradoxes
- Hilbert space is just a geometrical space with special features, such as:
  - It has asymmetric spiky shape (due to probability constraints)
  - Its true representation is a generalised Bloch body
  - It is measured using Fubini-Study fidelity metric
  - Its pure state manifold in highly crumpled within the outer Bloch boundary
- Therefore, all paradoxes translate to the projective Hilbert space manifolds, i.e.
  - Volume of the state manifold shrinks with dimensions
  - Volume of equatorial area expands exponentially (as a % of total volume)
  - Volume of pure states of zero entropy also vanishes exponentially
  - The interior of mixed states expands exponentially into high entropy hell
  - Cost landscape flattens as gradients approach zero (barren plateaus)
  - Random initialisation makes optimisation more difficult
  - All states become equidistant and orthogonal (orthonormal desert)
  - The path towards the optimum solution ceases to exist
  - Global measurements lead to sparsity of outcomes and imprecision
- This makes high-dimensional Hilbert space an amazing place!
- This also makes development of large quantum models very difficult!

Jacob's recent projects in collaboration with Sebastian Zając and associate researchers:

Quantum Information Field Theory (QIFT)
combining Quantum Field Theory with Information Processing

Geometry of the Hilbert space
in hybrid quantum-classical models

Quantum time series analysis
processing of quantum time series and signals, e.g. data encoding, measurement, and model training (such as QNNs, QAEs, QRC, etc.)

Quantum model expressivity and trainability
balancing the model's ability to effectively represent data in the Hilbert space, as well as, its capacity to learn and generalize

Impact of barren plateaus countermeasures on QNN capacity to learn
incl. methods of measuring its capacity

Quantum graph theory
exploration of quantum representation and processing of graphs e.g. identifying clusters and communities

Foundations of quantum machine learning
Business value of quantum computing and QML

Available resources:

https://jacobcybulski.com/
https://github.com/ironfrown/

https://github.com/ironfrown/pennylane_intro_hands_on/

Jacob Cybulski,
Founder, Enquanted, Australia
Hon. Assoc. Prof. Deakin, SIT

Parameterized Quantum Circuits (PQC) & Variational Quantum Algorithms (VQA)

Appendix

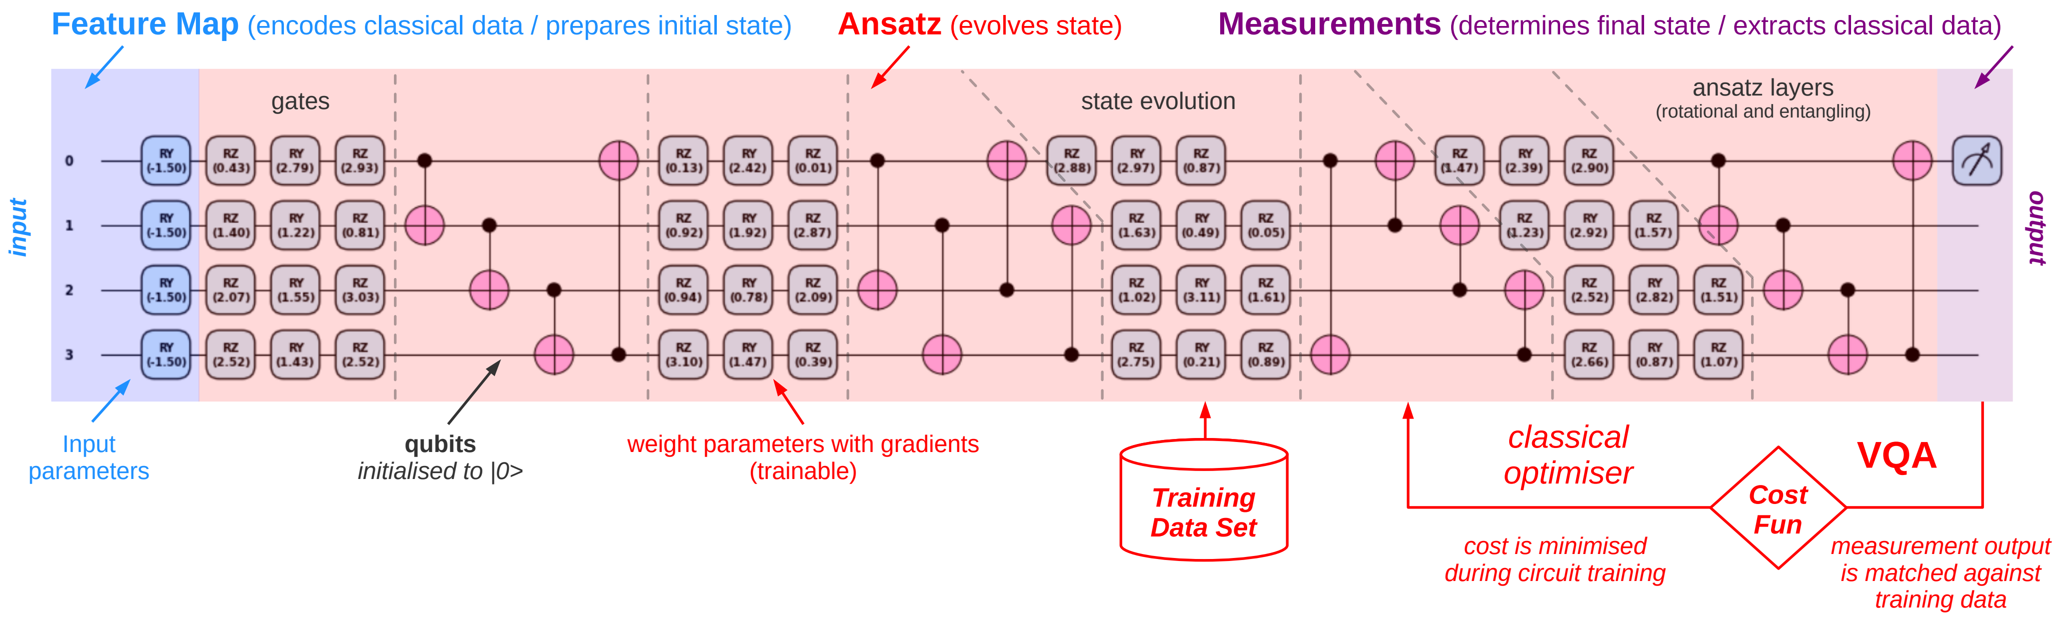

We create a **PQC** (variational model) = a circuit template with parameterised gates, e.g. Rx(a), Ry(a) or Rz(a) or P(a), rotating qubit state in X, Y or Z axis.

What does it mean a PQC is highly dimensional?
Lots of qubits. Lots of states. Lots of parameters.
But also… Lots of entanglements. Lots of measurements.