Predicting antibiotic resistance with quantum projective learning

Antibiotic resistance prediction is a hard ML problem for **structural reasons**, not just scale:

- High-dimensional, heterogeneous clinical features
- Severe class imbalance for most antibiotics
- Strong feature redundancy
- Decision boundaries that vary across antibiotics

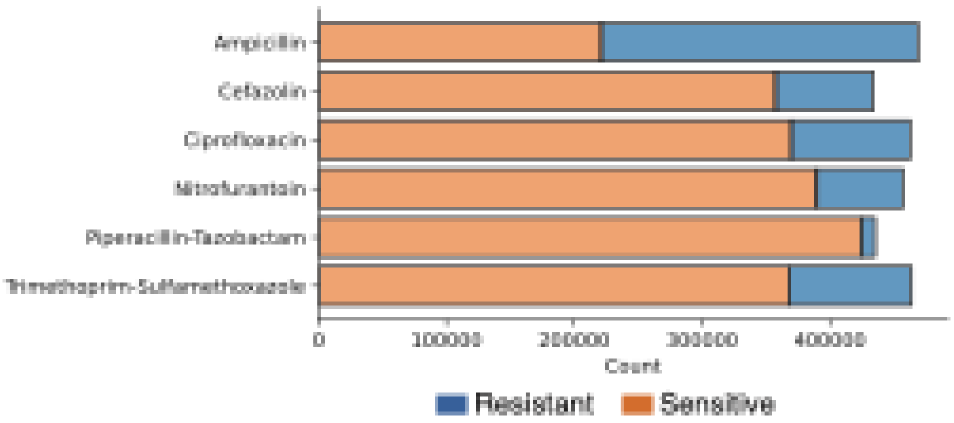

From a mathematical standpoint, this is a regime where **representation can matter more than model choice**.

Saw consistent performance across a range of classical MLs with AUCs ranging 0.71-0.90

Are there data regimes where a **quantum-defined feature map** produces a representation with measurably different geometry than classical embeddings?

This study is probing **representation geometry**, not only leaderboard performance.

Clinical data → geometry

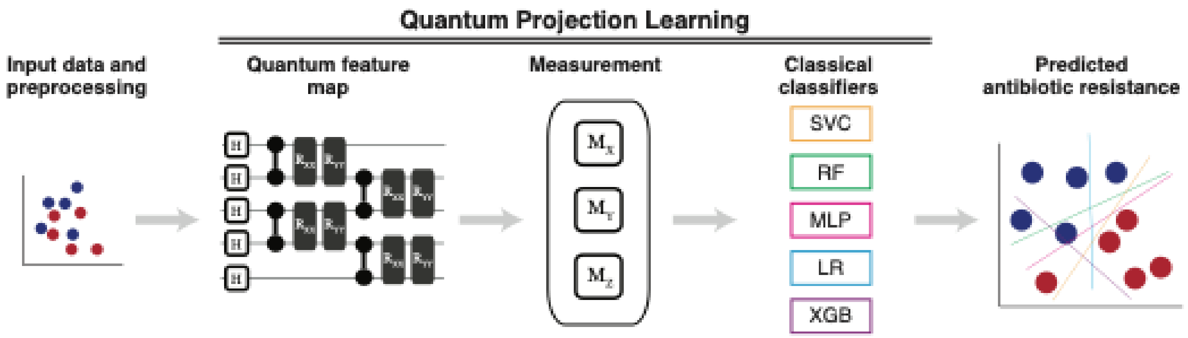

1

Strip away the clinical context and the problem becomes:

$$\{(x_i, y_i)\}_{i=1}^{N}, x_i \in \mathbb{R}^d,\ \ y_i \in \{-1, +1\}$$

2

Goal: Learn a classifier whose **generalization depends on the geometry** induced by a feature map

$$\Phi: \mathbb{R}^d \to \mathcal{H}$$

3

In **classical ML**, Φ might be:

- linear (PCA),
- kernel-induced (RBF),
- or implicit (trees, boosting)

4

In **Quantum Projection Learning**, Φ is *explicitly defined* by physics:

$$x \ \mapsto \ |\,\psi(x)\rangle = U(x)\,|\,0\rangle$$

followed by **projection**:

$$\Phi_q(x) = \left\{\left\langle \psi(x) \,\middle|\, \sigma_k^{(i)} \,\middle|\, \psi(x) \right\rangle\right\}$$

This is not a black box network; it is a **map into expectation values of observables**.

5

Despite partial theory on quantum kernels and projections, we don't know:

- When real, messy data lands in those favorable regimes
- How noise, depth limits, etc. collapse expressivity
- Does observed separation survive projection to low-order marginals (1-RDMs)

6

Given all practical constraints, when does the geometry induced by a quantum map differ enough from classical maps to change learning behavior?

That cannot be answered by asymptotic theory alone.

Evaluated 6 antibiotics in simulation and on QPU

While QPL does not outperform classical baselines in aggregate, there are datasets where QPL is finding some benefit

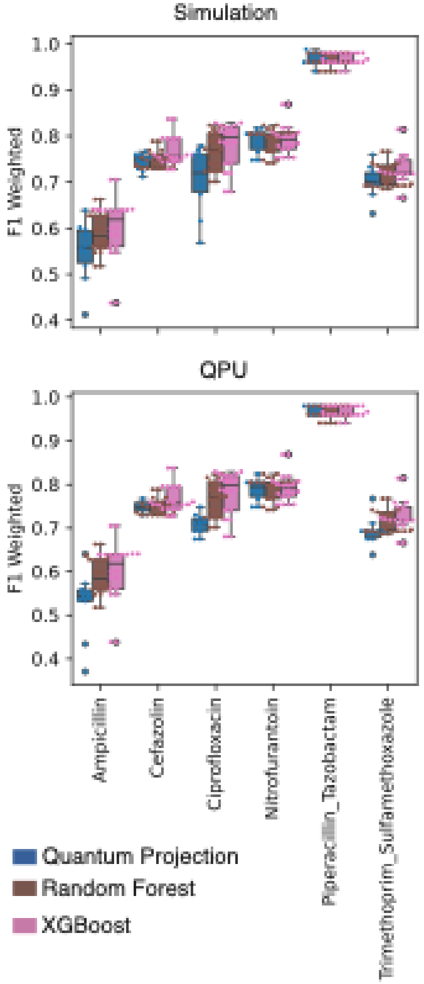

Performance per experiment

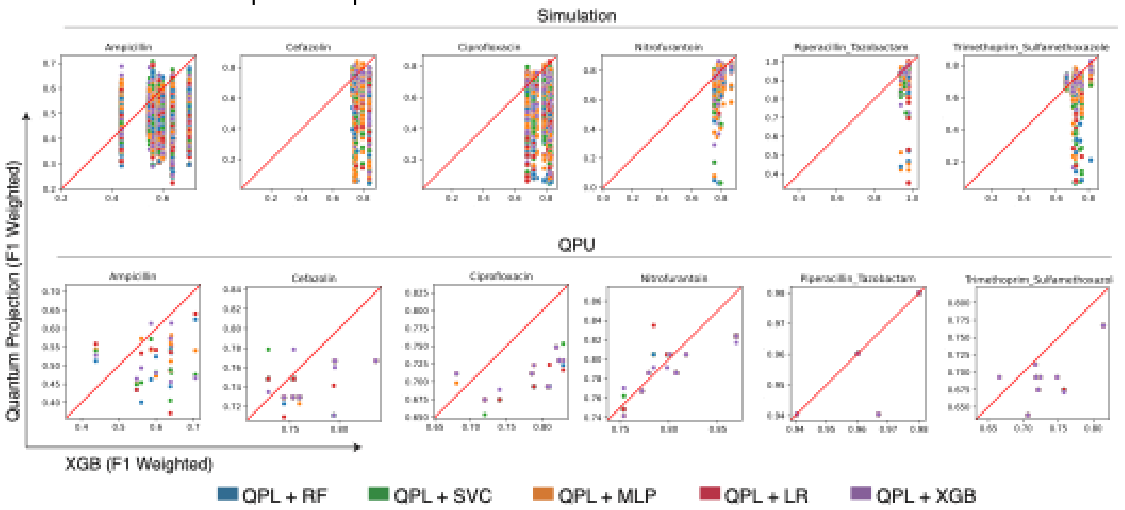

Complexity measures

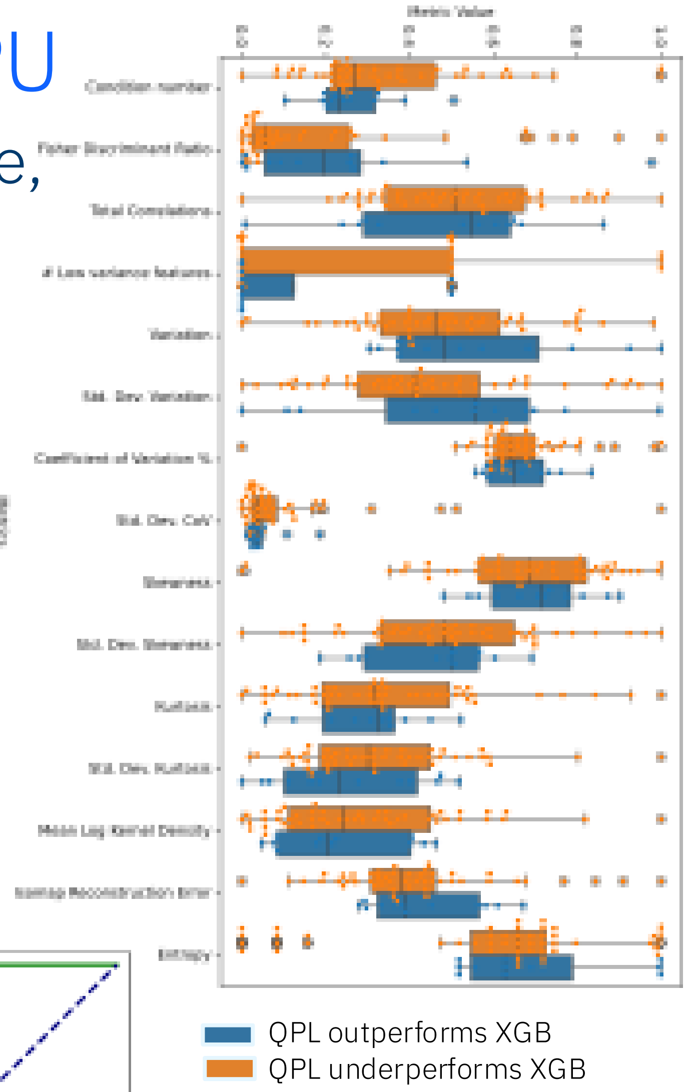

We identify a complexity measure signature associated to QPL performance

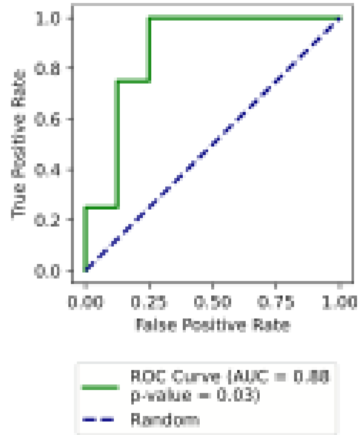

Conditional utility is still utility

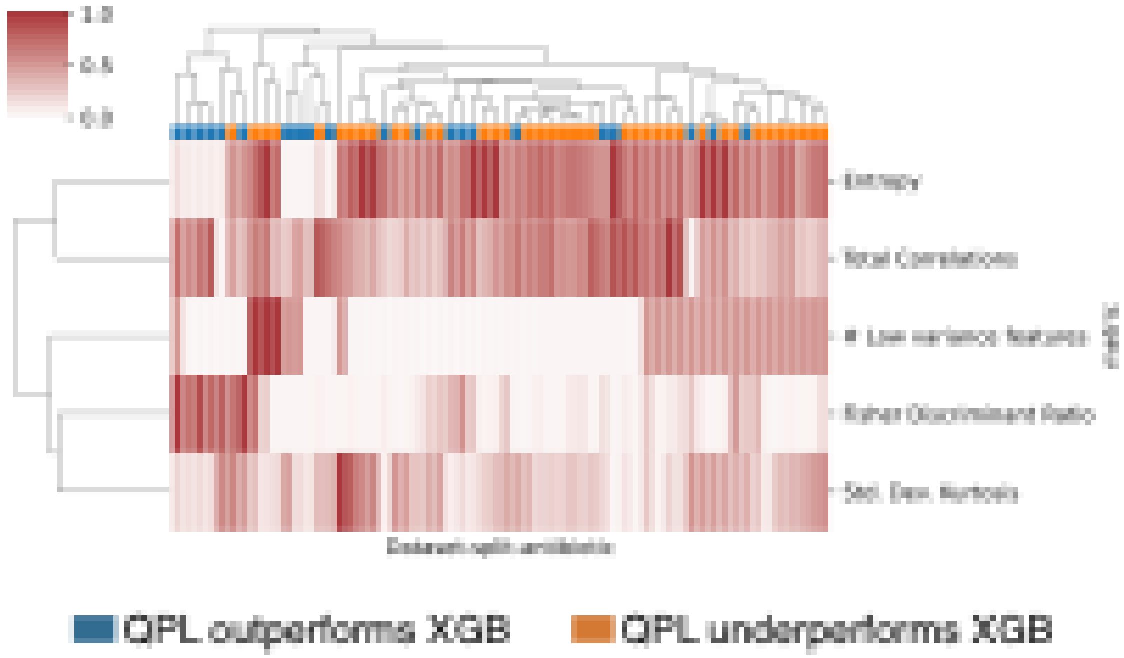

Here, QPL **did not surpass** classical baselines overall

Yet a **multivariate data-complexity signature** predicts when QPL *will* outperform (AUC = 0.88)

This tells us:

- Quantum projections are sensitive to certain features in the data that we can detect, such correlation structure or entropy

The question becomes:
“Given a dataset, can we decide whether a quantum embedding is worth invoking?”

That is a workflow-selection problem using as factors:

- Data complexity
- Kernel geometry[1]
- Model complexity[1]

Open unknowns:

- How expressive are low-depth, low-measurement quantum maps in practice?
- When does dimensionality reduction destroy quantum-useful structure?
- Can data complexity metrics predict kernel separation a priori?

1. Huang et al. Nat. Comms. 2021

Focusing on QPU experiments, we identify a complexity measure signature associated to QPL performance

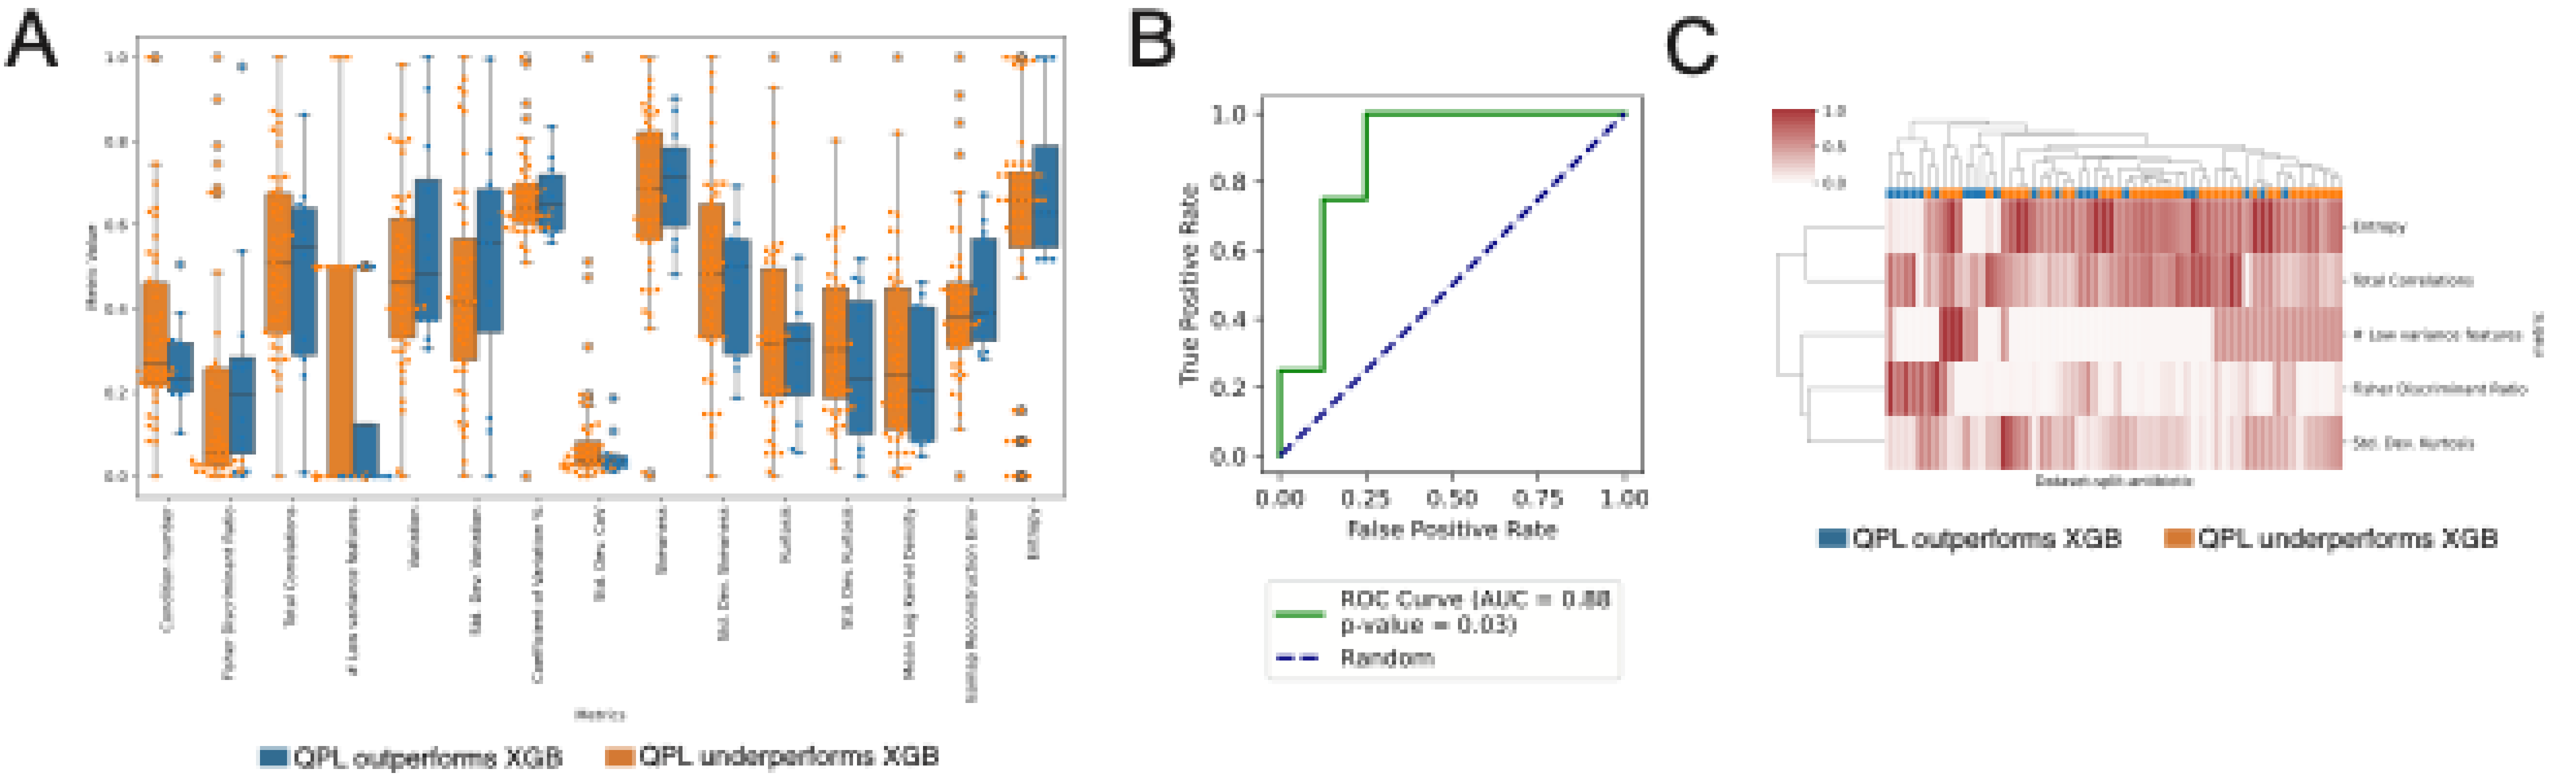

Conditional utility is still utility

Here, QPL **did not surpass** classical baselines overall

Yet a **multivariate data-complexity signature** predicts when QPL *will* outperform (AUC = 0.88)

Top features include:

- Shannon entropy
- Fisher Discriminant Ratio
- Total correlations
- Kurtosis variability
- Feature redundancy

This tells us:

- Quantum projections are sensitive to certain features in the data that we can detect, such correlation structure or entropy
- These are regimes where kernel approximations can struggle classically

The question becomes:
“Given a dataset, can we decide whether a quantum embedding is worth invoking?”

That is an optimization and workflow-selection problem using as factors:

- Data complexity
- Kernel geometry[1]
- Model complexity[1]

Open unknowns:

- How expressive are low-depth, low-measurement quantum maps in practice?
- When does dimensionality reduction destroy quantum-useful structure?
- Can data complexity metrics predict kernel separation a priori?

Not new models — **better understanding of representations.**

1. Huang et al. Nat. Comms. 2021