Breast Tumor Classification Using Quantum Neural Networks

Matt Purkeypile
August 2025
AAI590 Capstone
Masters of Applied Artificial Intelligence
University of San Diego

Agenda

- Overview of data set
- Background:
  - Quantum computing
  - Quantum neural networks
- Experimental methods:
  - Static
  - Adjusted
- Optimizations
- Results
- Conclusion
- Questions and resources

Objective

The hope is that this work serves as an inspiration and example of the immense potential of the blending of quantum computing and artificial intelligence.

Overview of data set

- Breast Cancer Wisconsin Diagnostic dataset from Kaggle.
  - 569 records.
  - Clean: no missing data or extreme outliers.
  - Describes different characteristics of tumors.
  - Skewed normal distributions.
- Data: identifier, 30 features, 1 binary target.
  - The 30 features consist of 10 features with 3 variations: mean, worst, and standard error.
  - Example: area occupied by the tumor. There is a mean area, worst area, and standard error of area.
  - Multiple features are highly correlated.

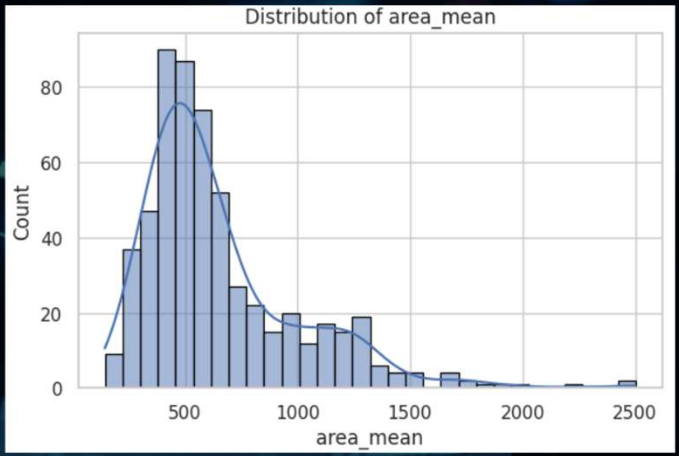

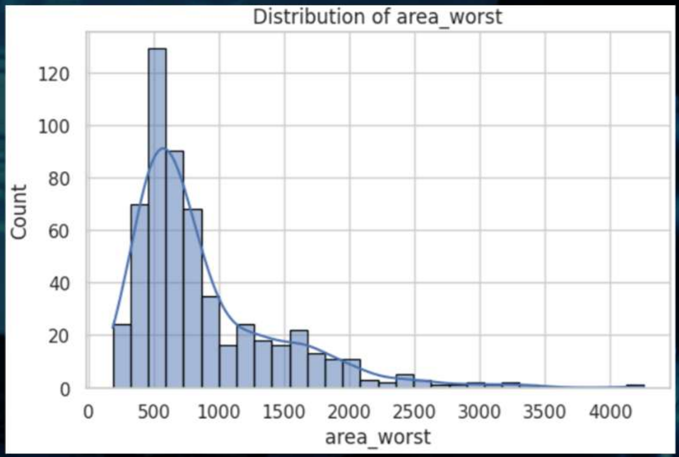

Overview of data set: correlation matrix

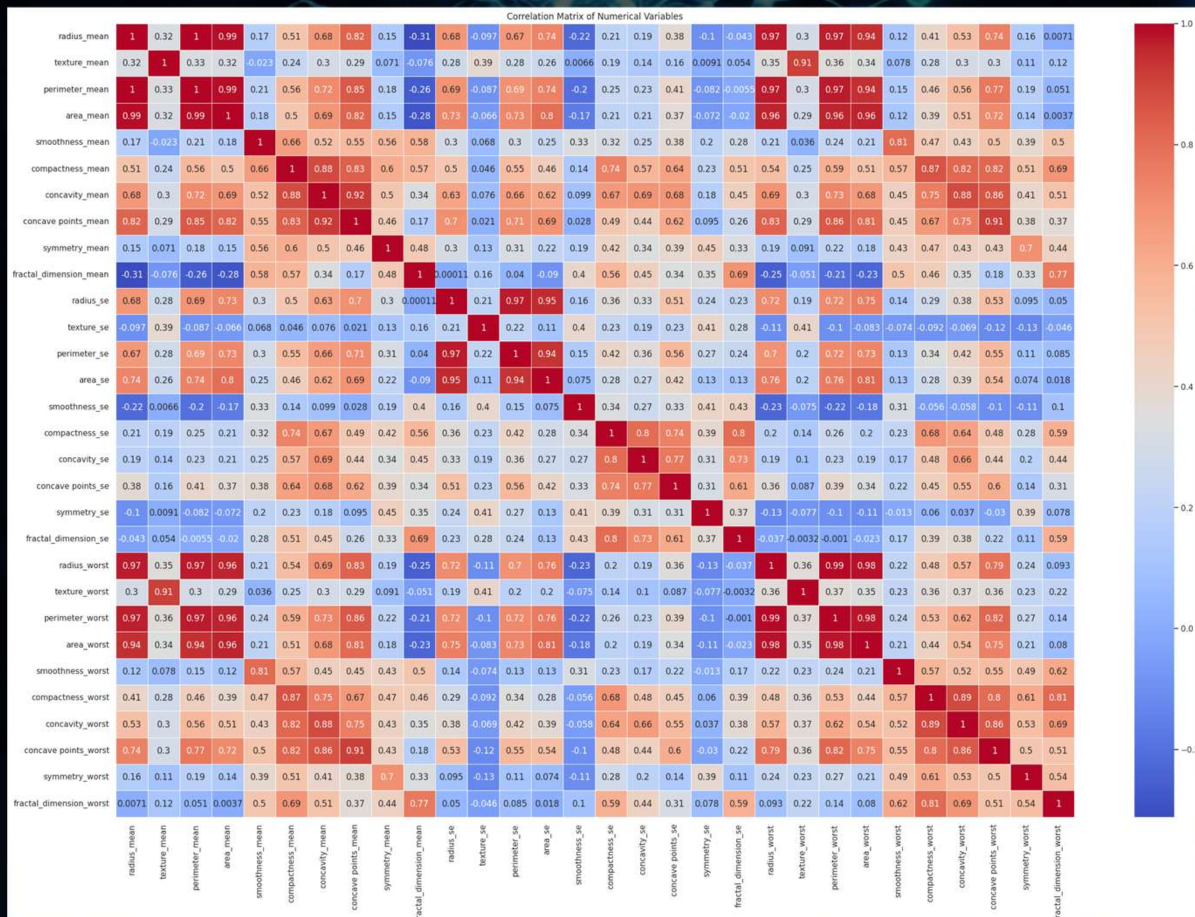

Background: quantum computing

- Existing computers, called classical, operate on bits: 0 or 1.
- Quantum computers operate on quantum bits, known as qubits. These can be a combination of 0 and 1 but exist in a higher dimensional space than bits or probabilistic bits.

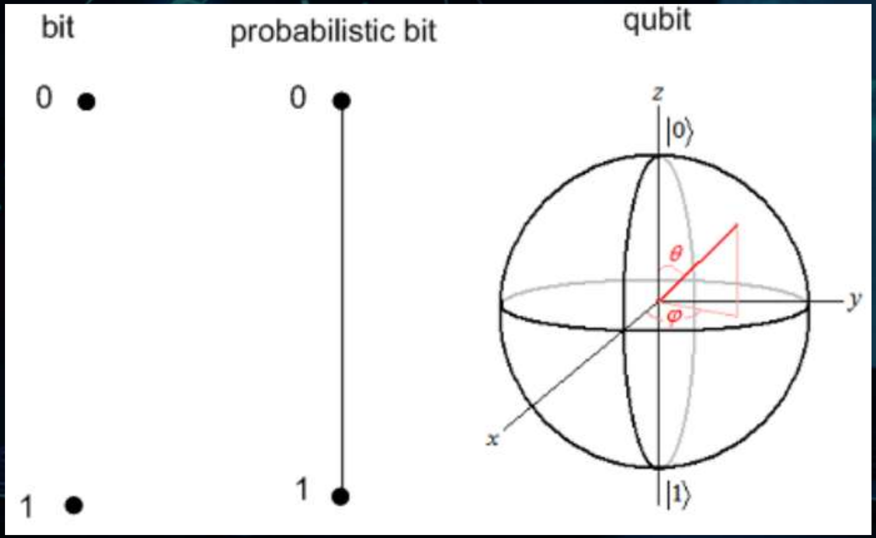

Background: quantum computing

- Unique aspects of quantum computers:
  - Superposition: the trait that a qubit can exist in a state that is a combination of 0 and 1.
  - Qubits are “measured”, which results in a classical 0 or 1 result.
  - Entanglement: tying the states of 2 or more qubits together.
    - Example: two qubits could be in a superposition of the states 00 and 11. Measurement will collapse those to either the state 00 or 11.
    - In this case, the result is non-deterministic.
  - Essentially allows for n qubits to represent 2^n states at once, while n bits can only represent one of those states at a time.
- Some restrictions: computations must be reversible, qubits cannot be copied and pasted (no cloning theorem).
- Myth: quantum computers are universally faster than classical computers. This is only true for certain problems.

Background: quantum computing

Example:

- Put a qubit in an equal super position of 0 and 1.
  - Mathematically expressed as $\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$
- Measure it to get a classical result of 0 or 1.
  - Since the qubit is in superposition, we essentially get a random 0 or 1 with equal odds.
- Repeat 20 times to get 20 random bits.

```
Circuit:
(0, 0): ───X^0.5───M('m')───
Results:
m=00110111000001011000
```

Background: quantum neural networks

- Similar to classical neural networks, except:
  - Classical features are mapped to quantum states.
  - The hidden layers are quantum circuits operating on qubits.
  - Measurement at the end results in classical output, usually non-deterministic.

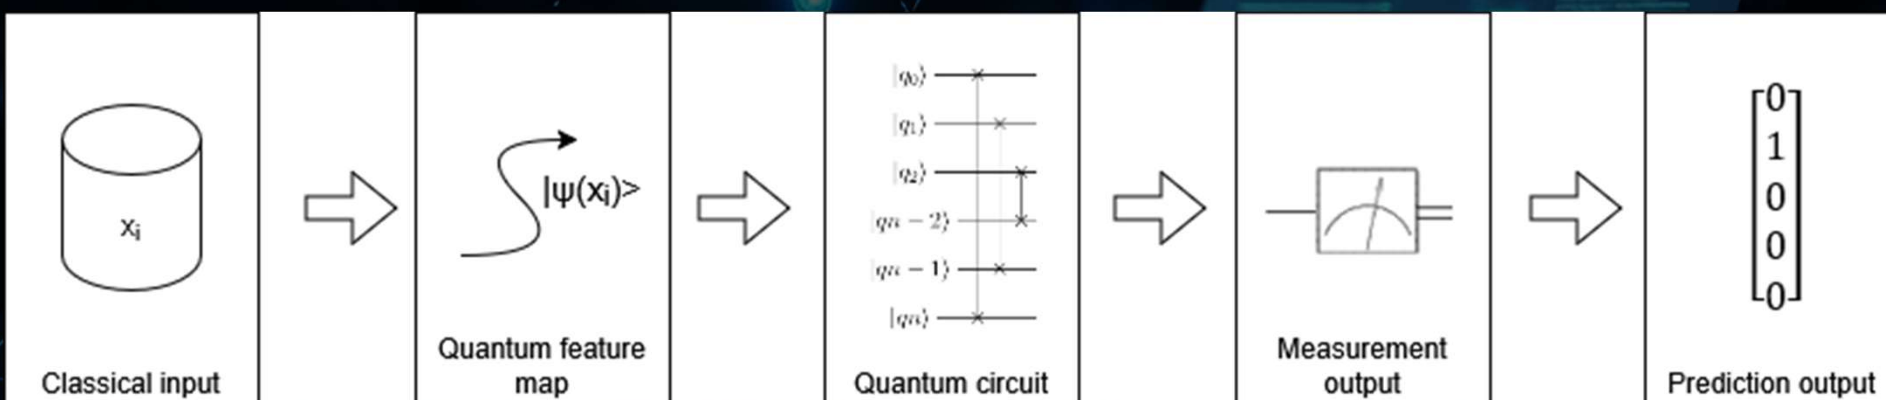

Experimental methods: static

The following were largely unchanged throughout the project:

- Use of hinge accuracy for the loss function.
- Each hidden layer using the same structure.
- Use of the Adam optimizer in TensorFlow, with default parameters.
- A standard data split of 80% training and 20% testing

Experimental methods: adjusted

Multiple components were adjusted to improve model performance:

- Number of epochs, although most attempts used 20 due to performance plateauing.
- Batch size: powers of 2 from 2^3 to 2^7 were attempted.
- Number of hidden layers: 1, 2, 5, and 8.
- Feature selection.
- Normalization techniques.
- Quantum feature mapping

Optimizations

- The components mentioned in the previous slide were adjusted.
- The best results (accuracy ~90%) were achieved with:
  - 20 epochs
  - Batch size of 64
  - Using the “worst” variation of 10 features (as opposed to the mean or standard error variations)
  - A quantum feature mapping of normalized data in the range of 0-1 of to:
    - $|0\rangle$when the classical value was <= 0.25
    - $\frac{1}{\sqrt{2}}|0\rangle+\frac{1}{\sqrt{2}}|1\rangle$ when the classical value as > 0.25 and <= 0.75
    - $|1\rangle$ when the classical value was > 0.75
  - 8 hidden layers

Results: base model performance

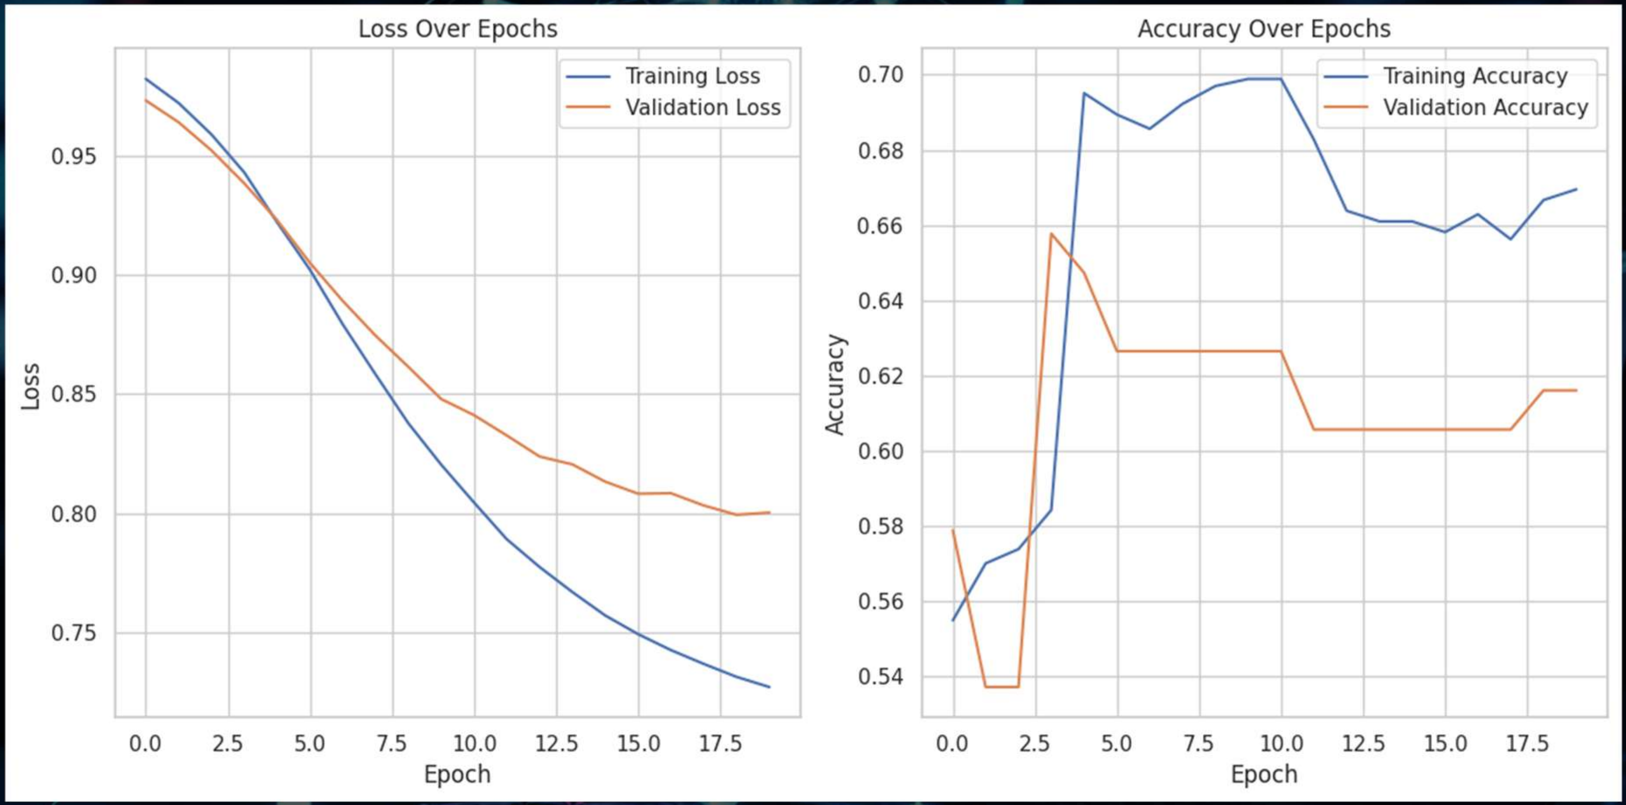

Final test loss: 0.8002 Final test accuracy: 0.6155

Results: optimized model performance

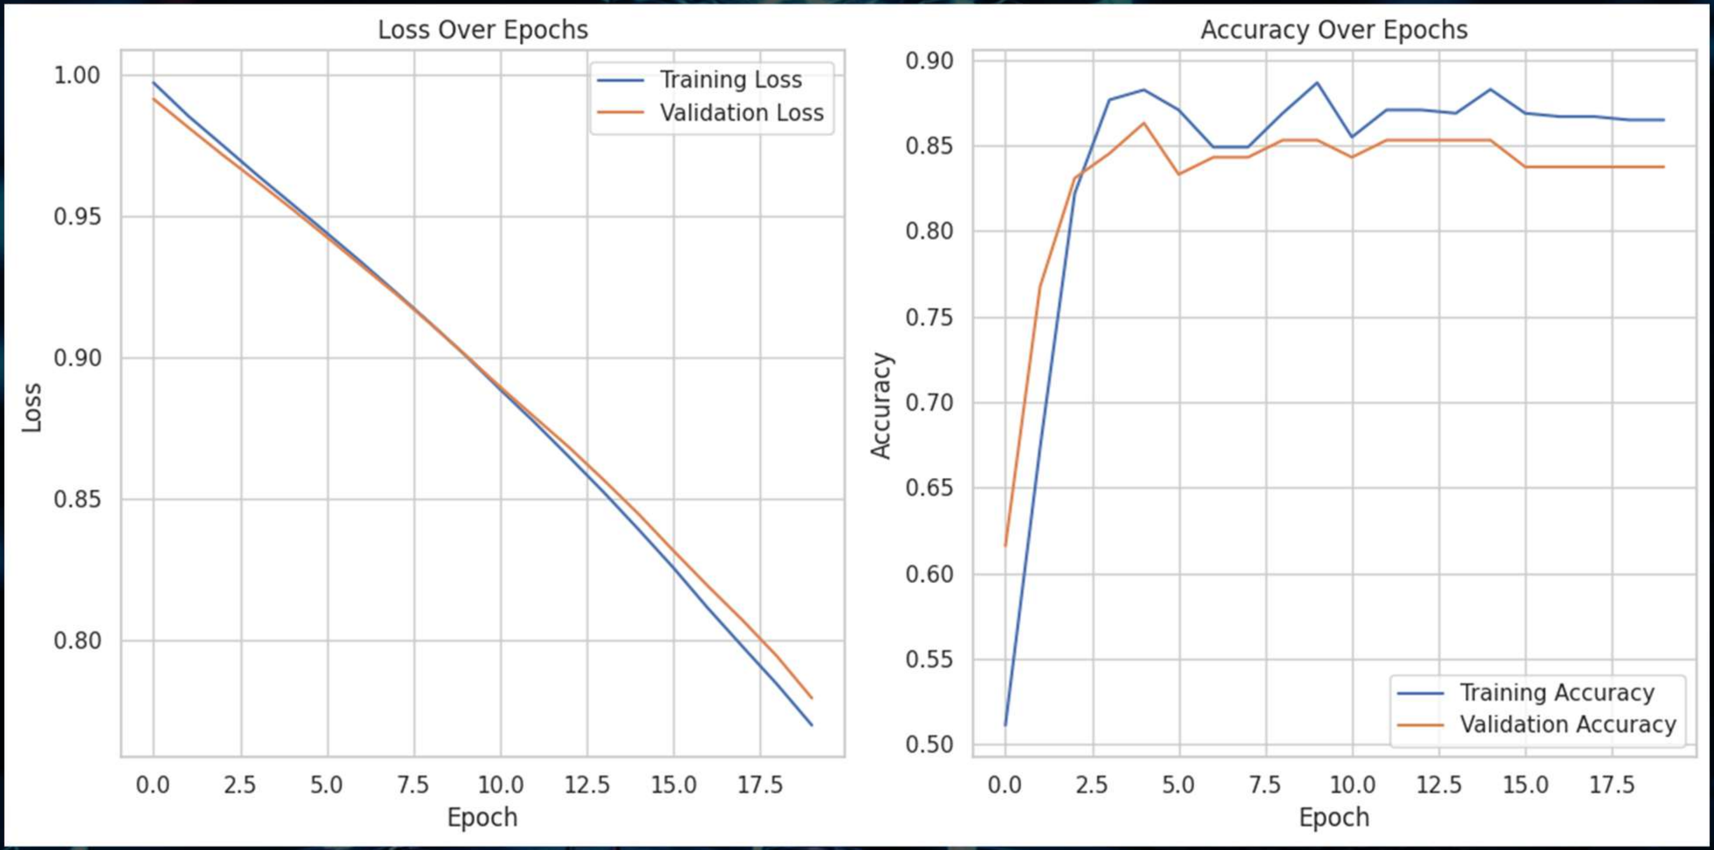

Final test loss: 0.9020 Final test accuracy: 0.9036

Results

- Able to detect cancerous tumors with ~90% accuracy
- This was done with only 10 qubits of information.
- Classical equivalent is 10 bits.
- The original Unicode characters required 16 bits to express a letter.
- This project demonstrates predicting cancer with high confidence using less information than it takes to store a single letter!

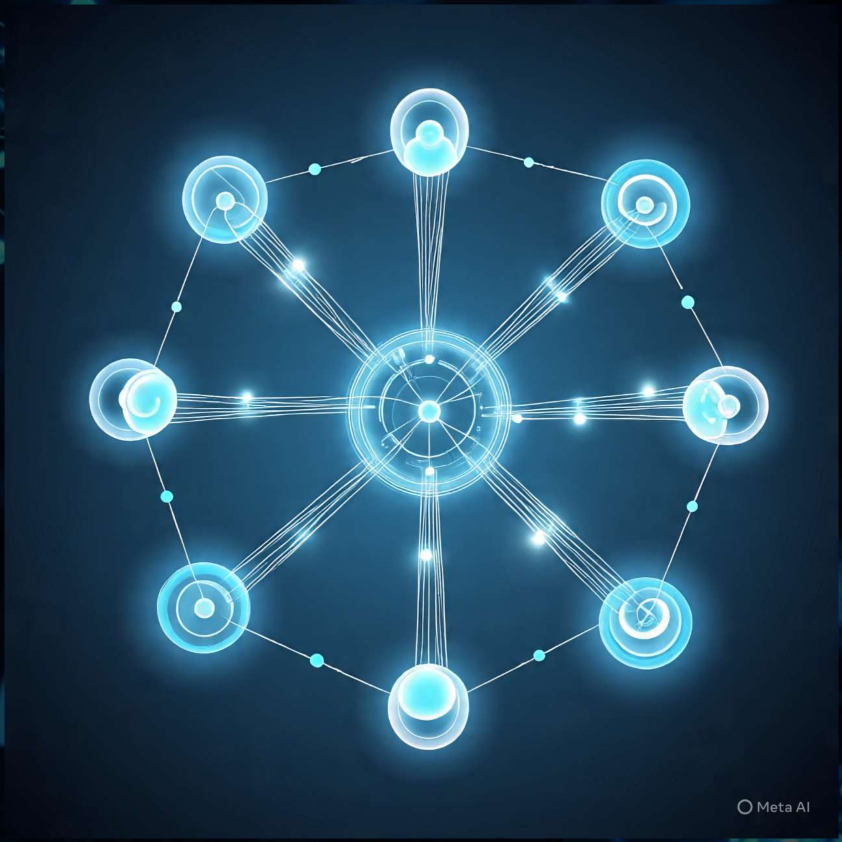

Conclusion

This project demonstrates the tremendous potential of leveraging quantum computing in machine learning. As Michio Kaku said:

Perhaps the key to artificial intelligence lies in the quantum theory. In fact, the merger of the two may revolutionize every branch of science, alter our lifestyle, and radically change the economy. AI will give us the ability to create learning machines that can begin to mimic human abilities, while quantum computers may provide the calculational power to finally create an intelligent machine.

Source: Kaku, M. (2023). Quantum Supremacy: How the Quantum Computer Revolution Will Change Everything. New York, NY: Random House LLC.

Questions?

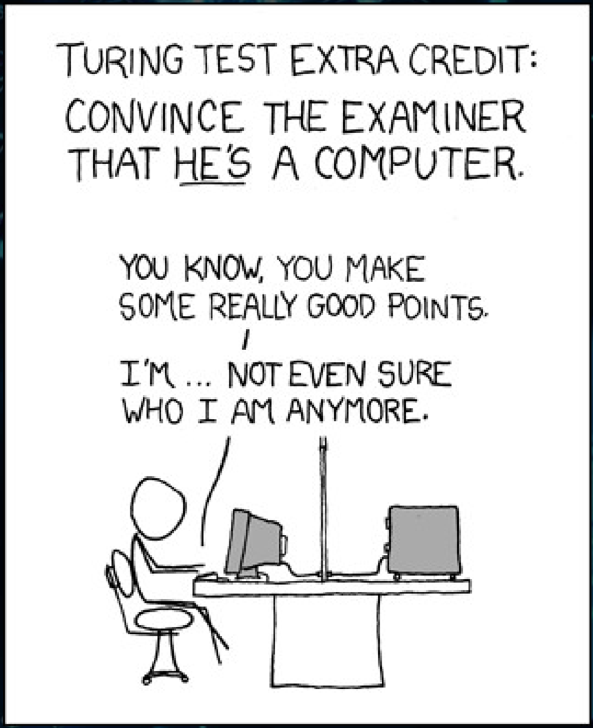

Image credit: XKCD, https://xkcd.com/329

Slide background created with Meta AI

Resources

- Code repository: https://github.com/mpurkeypile-usd/AAI590-Final_Project
- Final paper: http://cove.purkeypile.com/AAI590-Final_Project-MPurkeypile.pdf
- This deck: http://cove.purkeypile.com/AAI590-Final_Project-Presentation-MPurkeypile.pdf
- Author contact info: m purkeypile at san diego dot edu