Theoretical Physics in an era of Machine Learning

Frontiers of Physics - News from the NBIA

Evert van Nieuwenburg
evert.vn@nbi.ku.dk

The Niels Bohr International Academy

Theoretical Physics in an era of Machine Learning

Frontiers of Physics - N

Please feel free to ask questions!

The Niels Bohr International Academy

Some words about me

Leiden University, Netherlands → ETH zürich → California Institute of Technology (Caltech)

[My Trajectory On Google Earth]

The content of this lecture

What is Machine Learning? An introduction by example

How can it help research in physics?

TiqTaqToe

Bonus: Quantum Games

Machine Learning is fun

Locomotion

deepmind.com/blog/article/producing-flexible-behaviours-simulated-environments

Machine Learning is fun

Multi-Agent Hide & Seek

https://openai.com/blog/emergent-tool-use/

Locomotion

deepmind.com/blog/article/producing-flexible-behaviours-simulated-environments

Machine Learning is everywhere

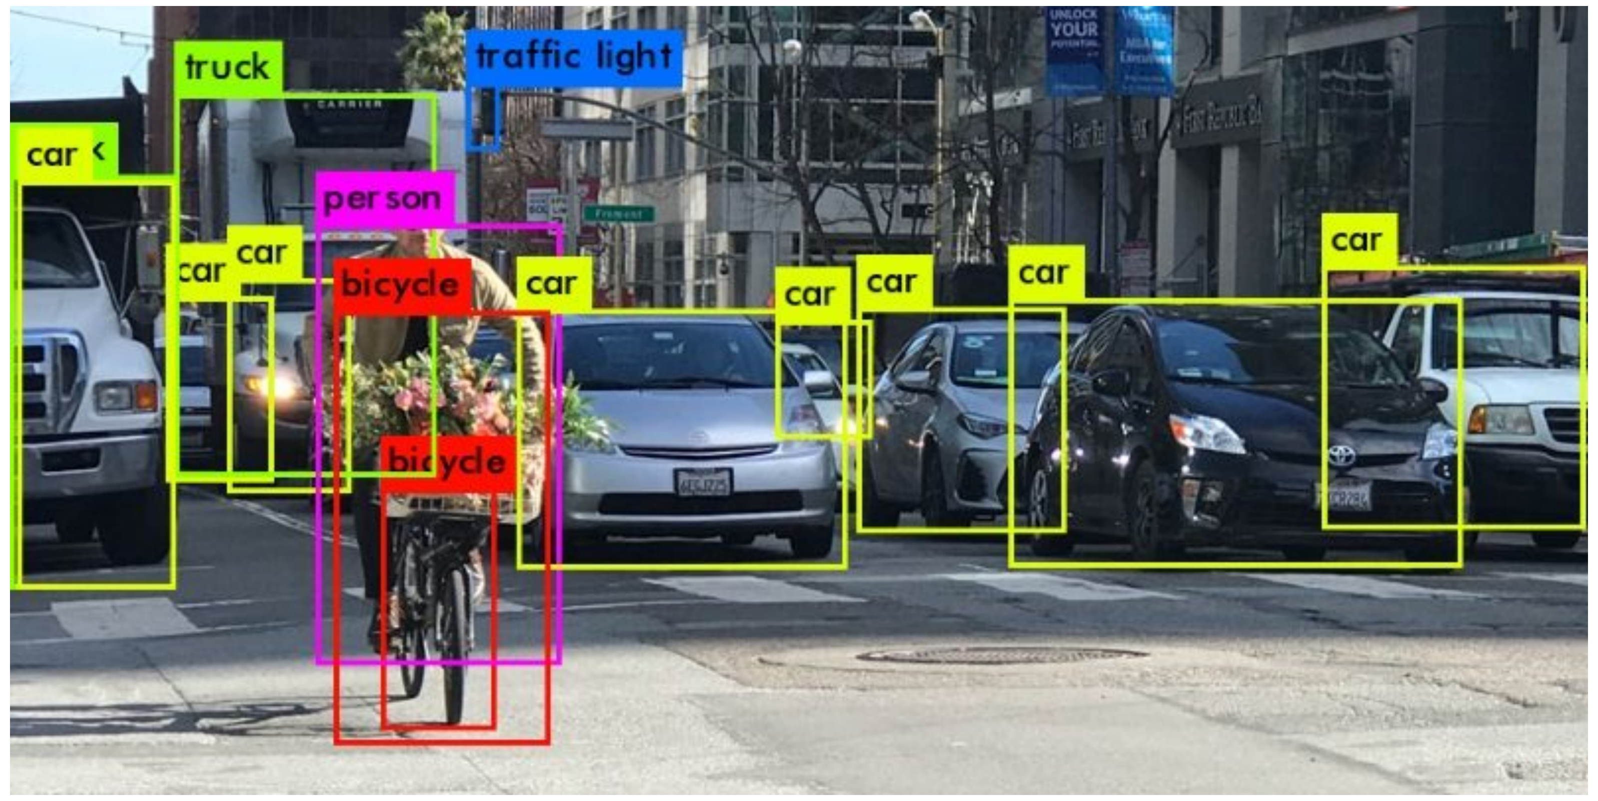

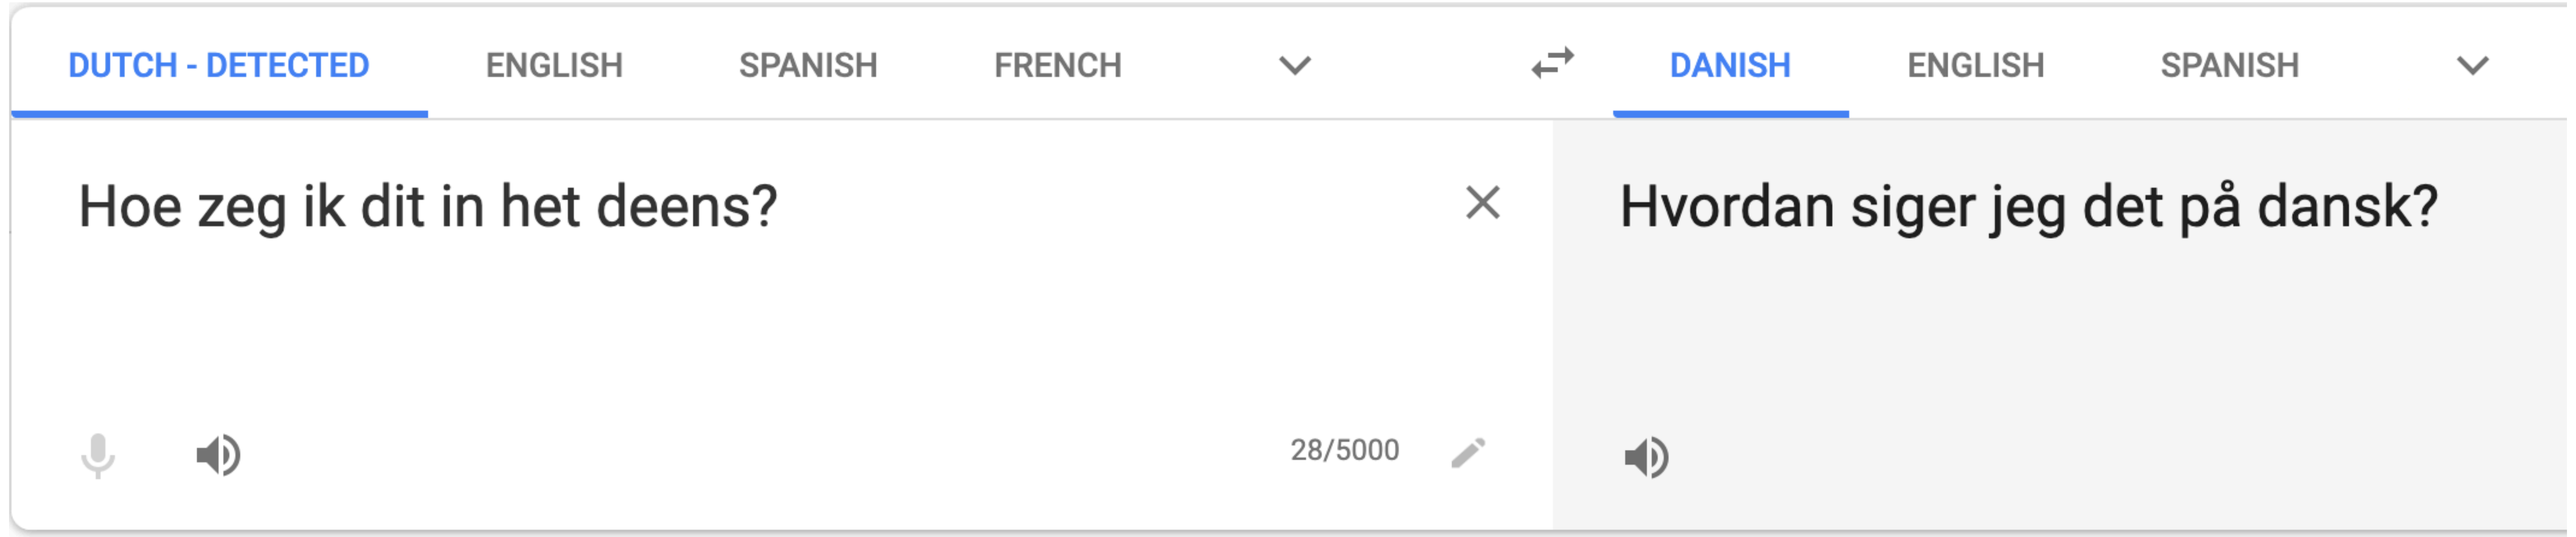

Natural Language Processing
https://talktotransformer.com/

Generative Modelling
https://thispersondoesnotexist.com/

Speech Synthesis
“Deep-Fakes”
…

Machine Learning is a different way of solving problems

(An example of "supervised learning" - more in a few slides!)

The ML Approach

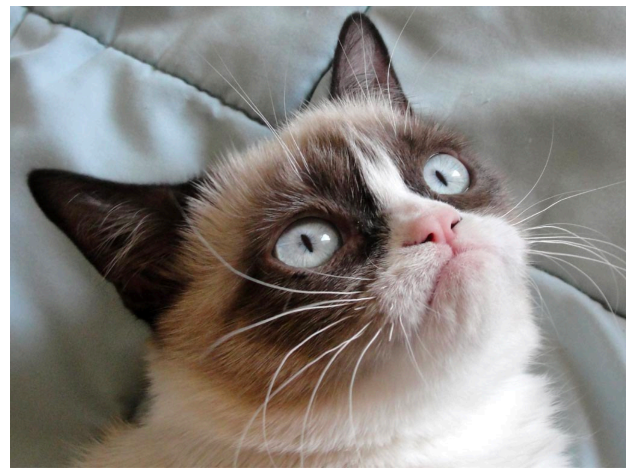

TRUE

FALSE

FALSE

TRUE

The Human Approach

```
FUNCTION IS_CAT?:
 IF( HAS_WHISKERS ):
   IF( HAS_CAT_EYES ):
     IF( HAS_NO_MANE ):
       IF( . . . ):
          RETURN TRUE
 ELSE:
   RETURN FALSE
```

The main component in many ML techniques is a neural network

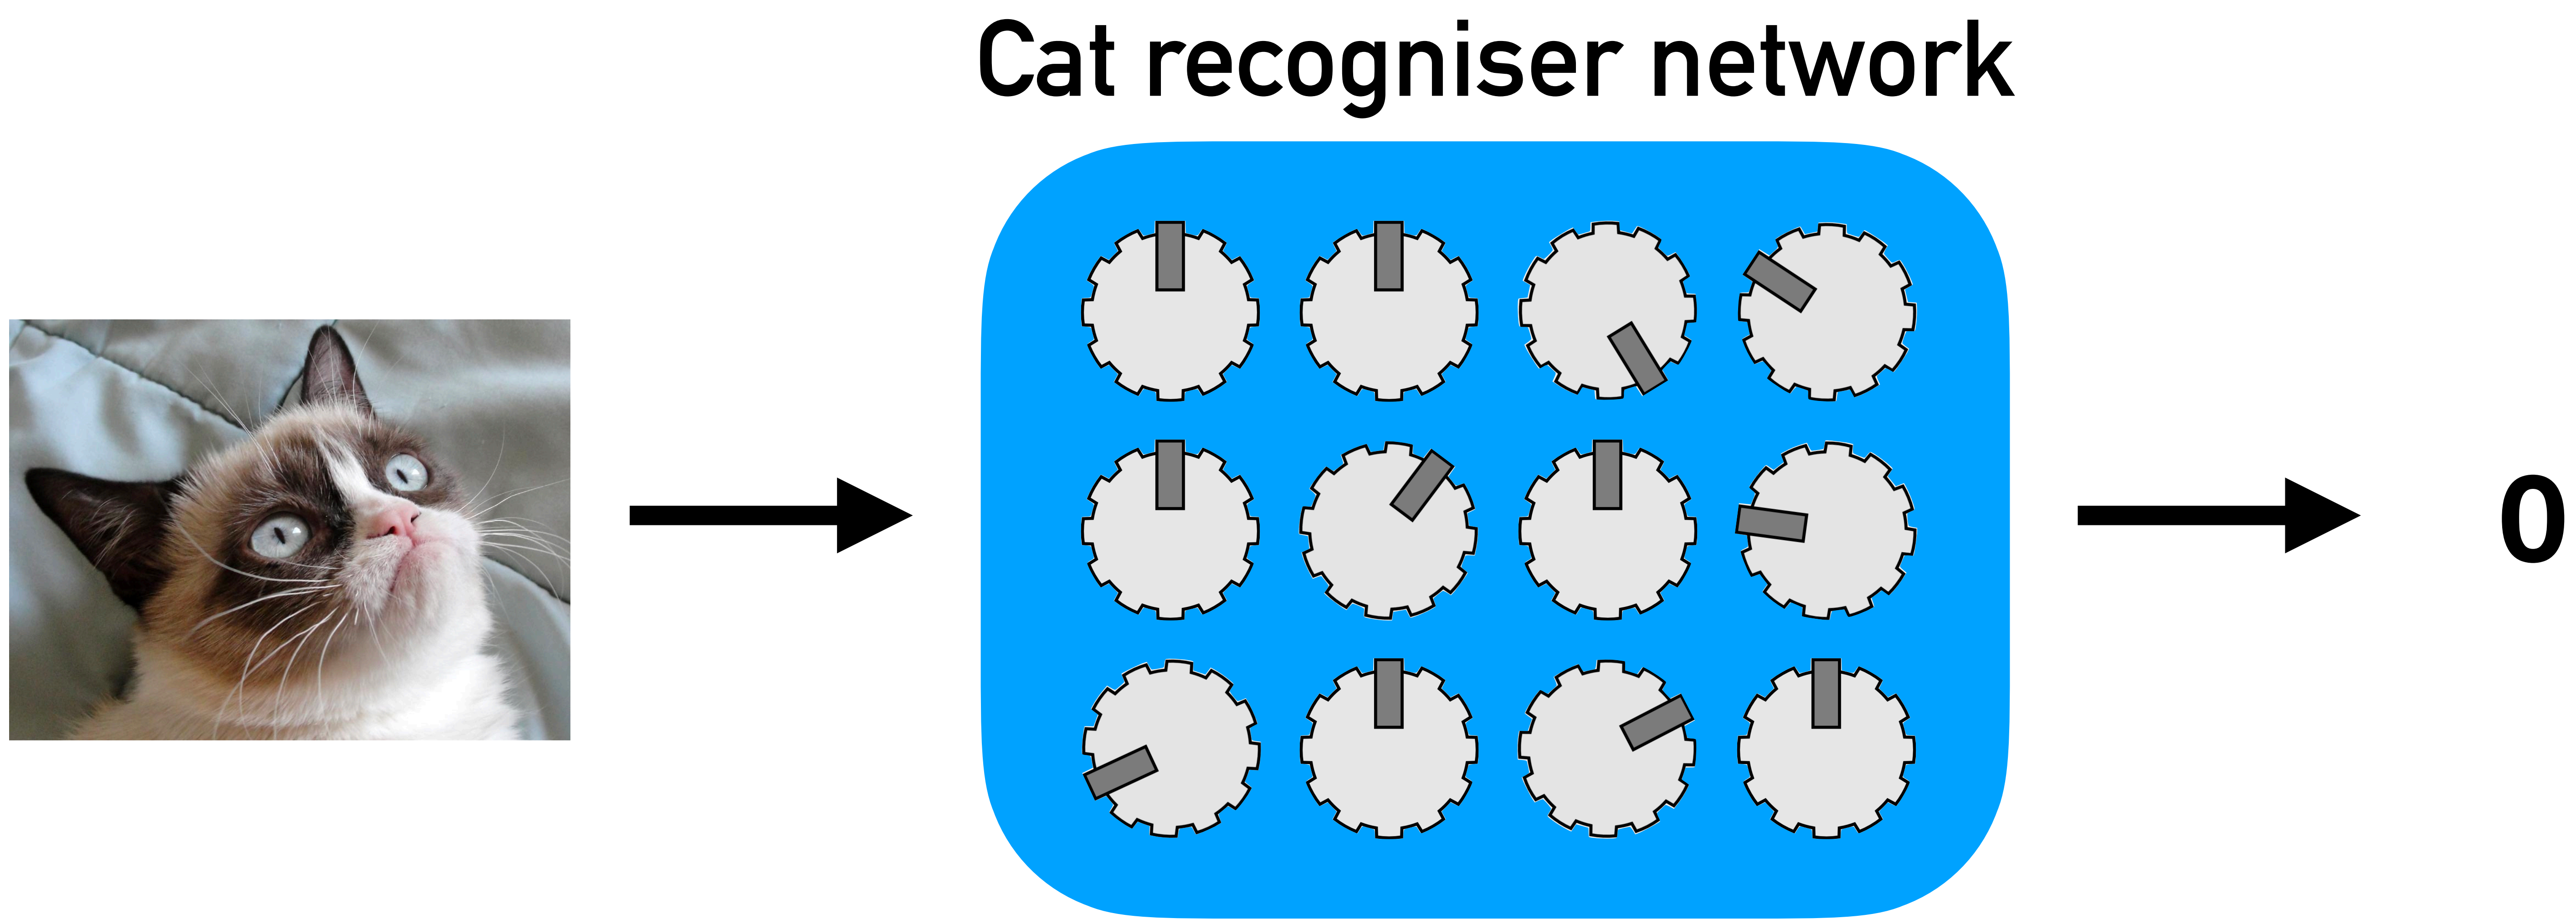

The main component in many ML techniques is a neural network

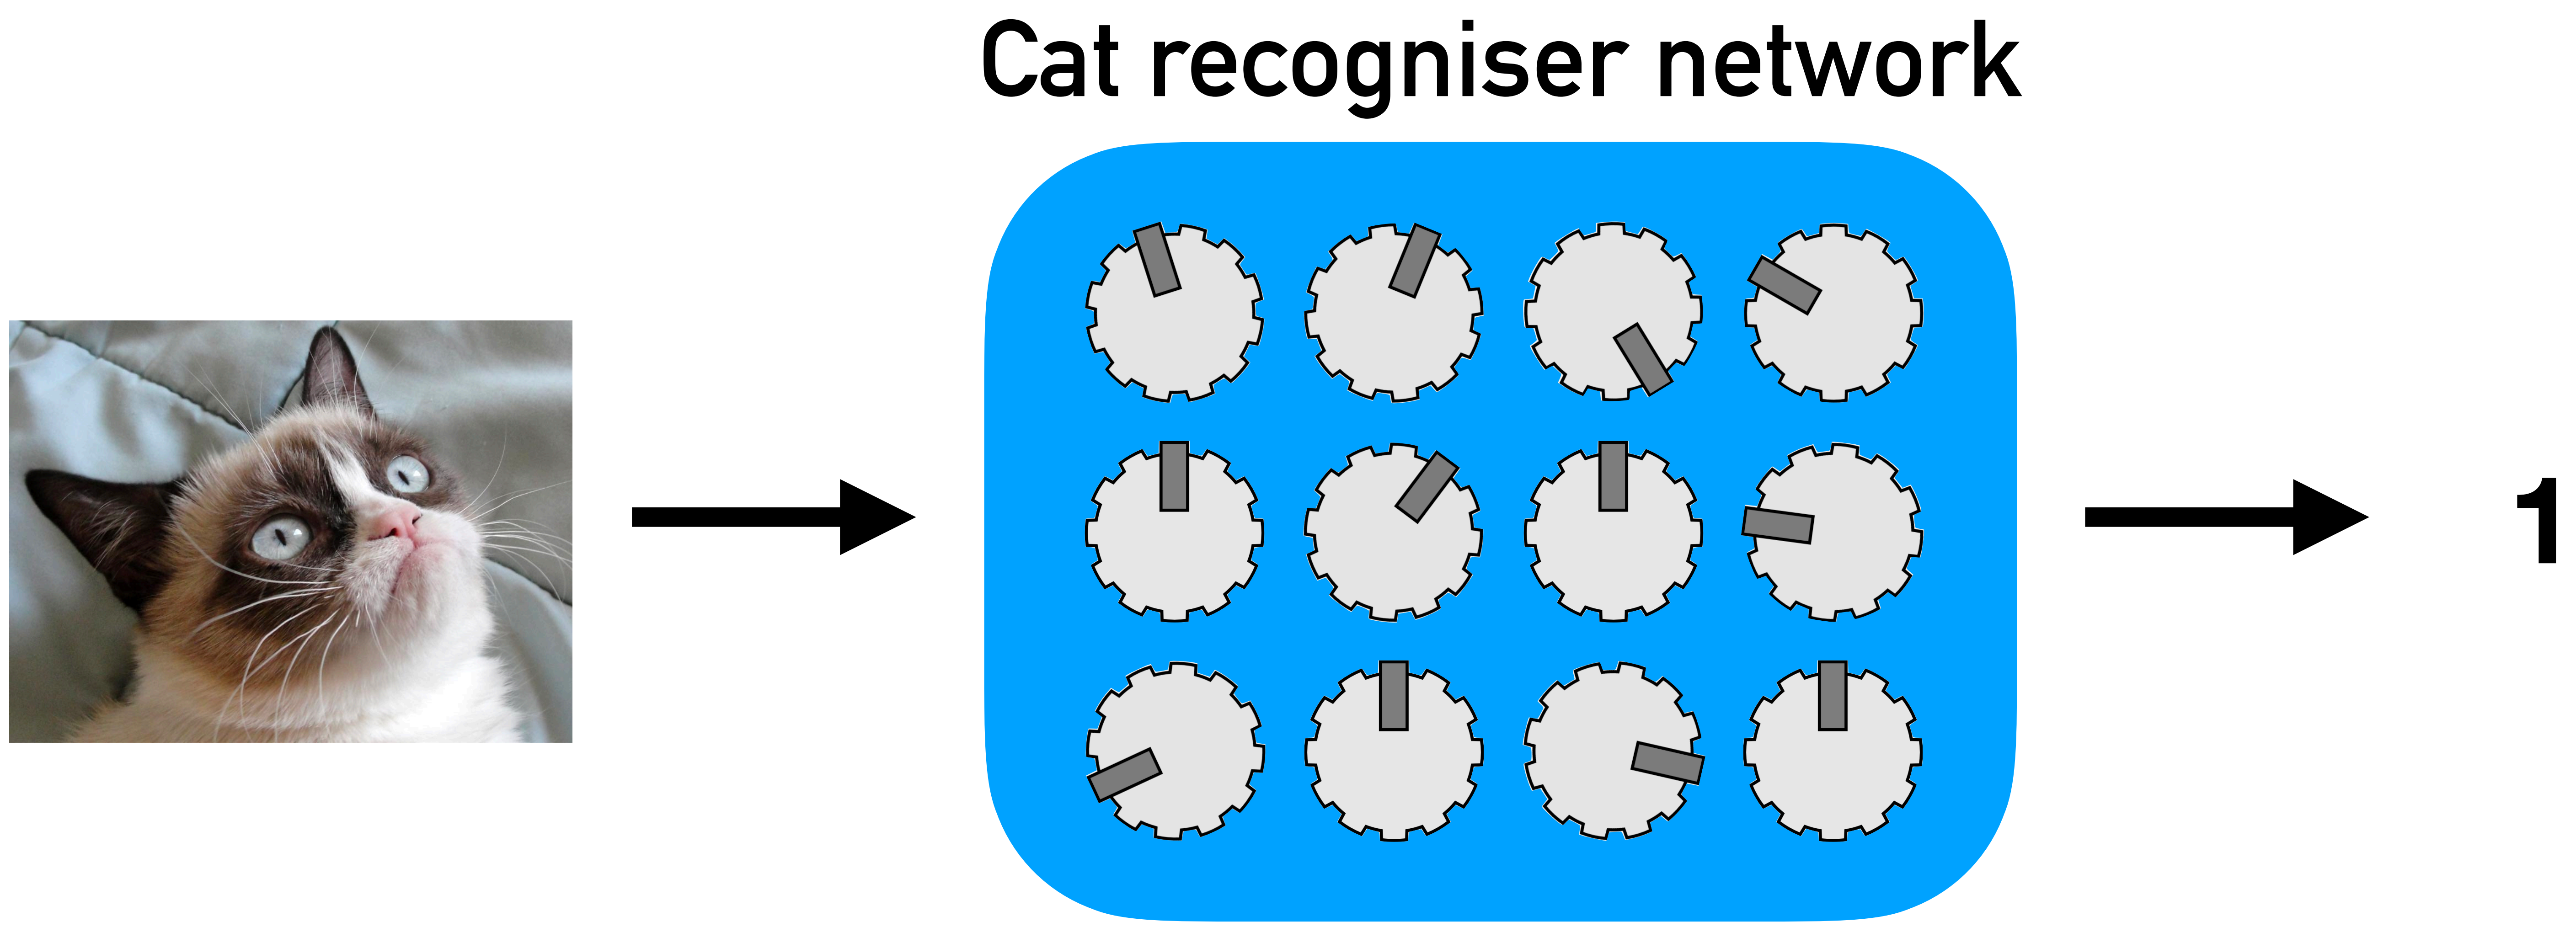

The main component in many ML techniques is a neural network

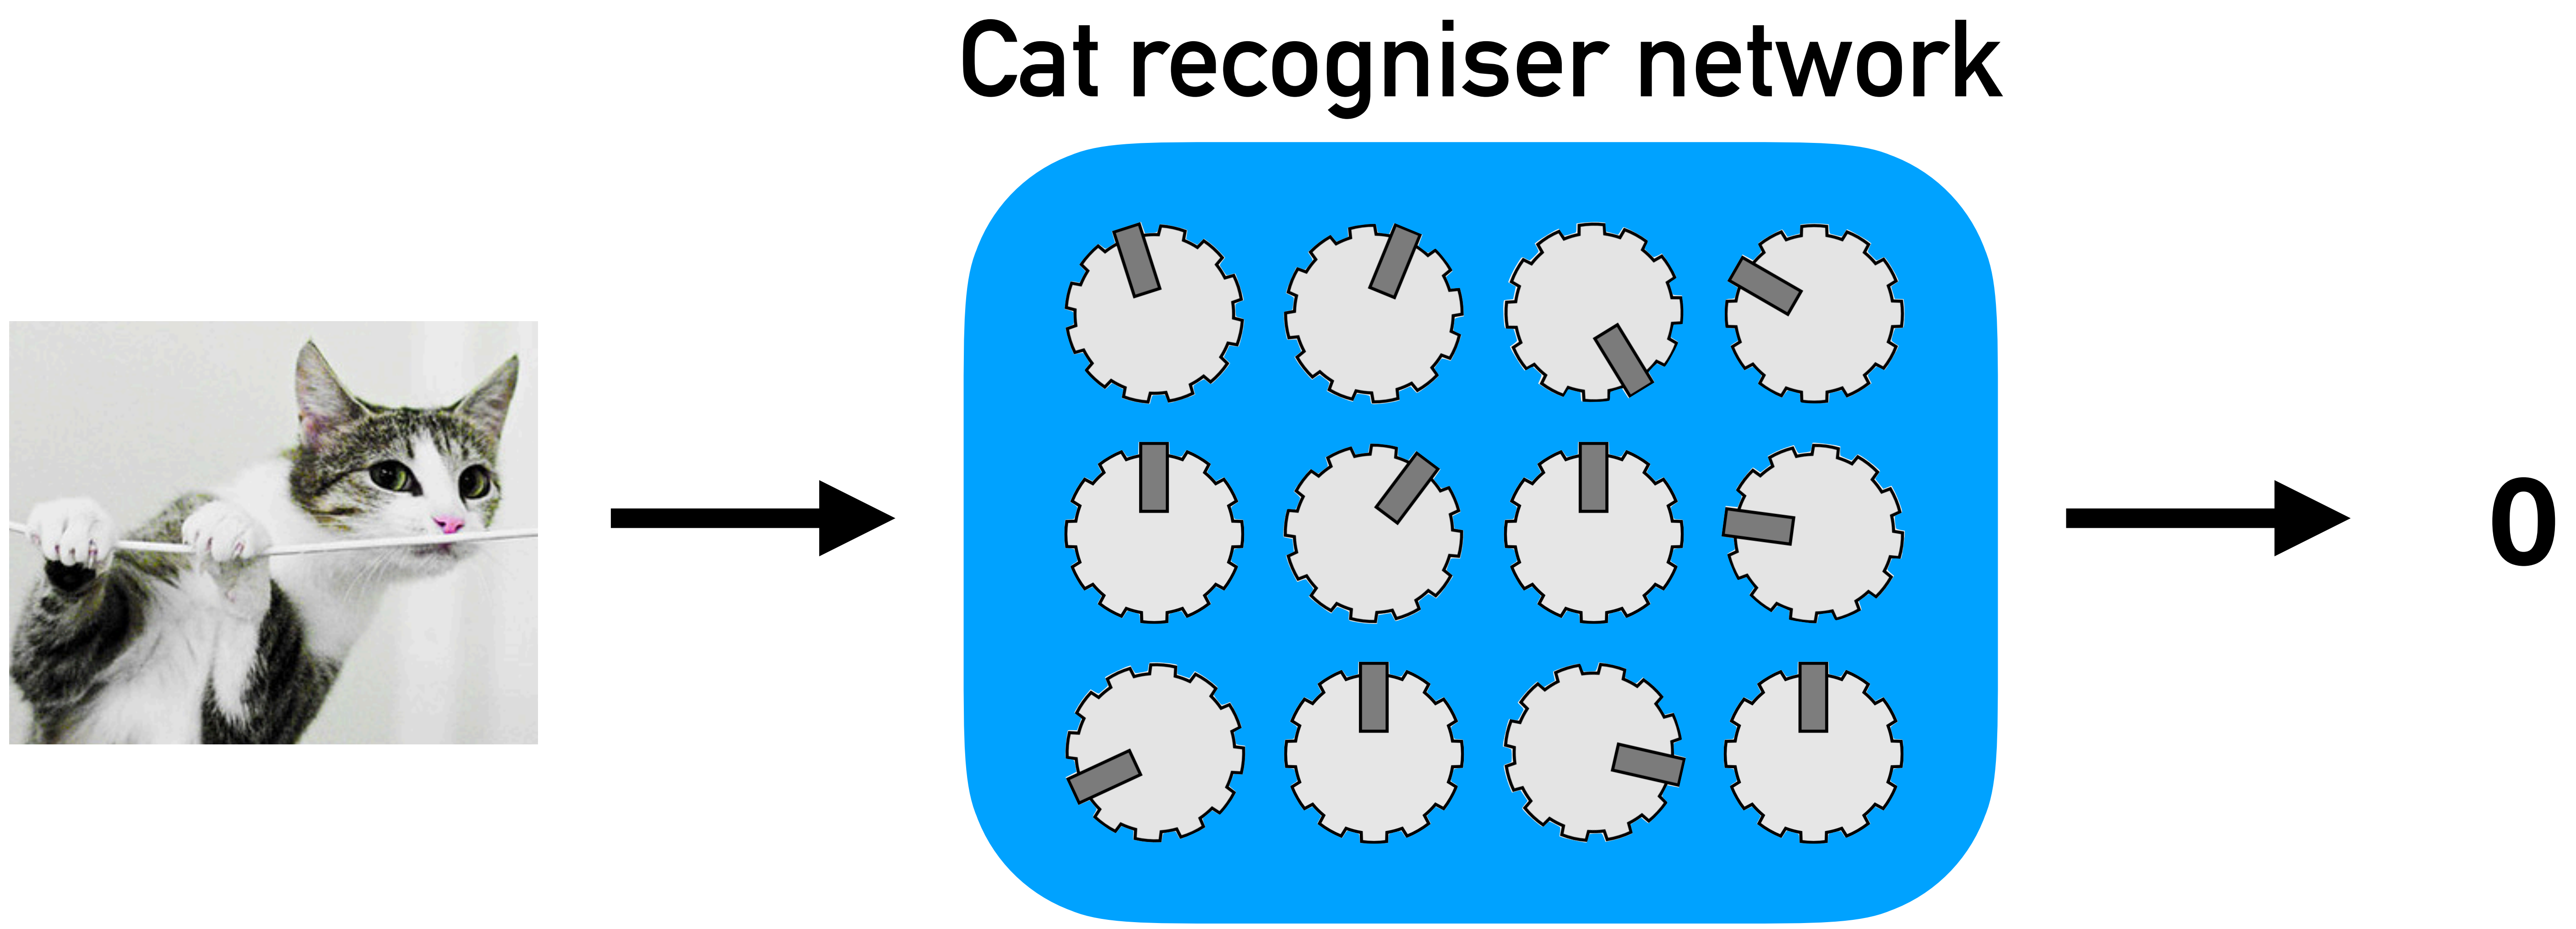

Repeat for all cats, tweaking the knobs so that we recognize all of them!

This is 'learning'

This is how a simple NN works

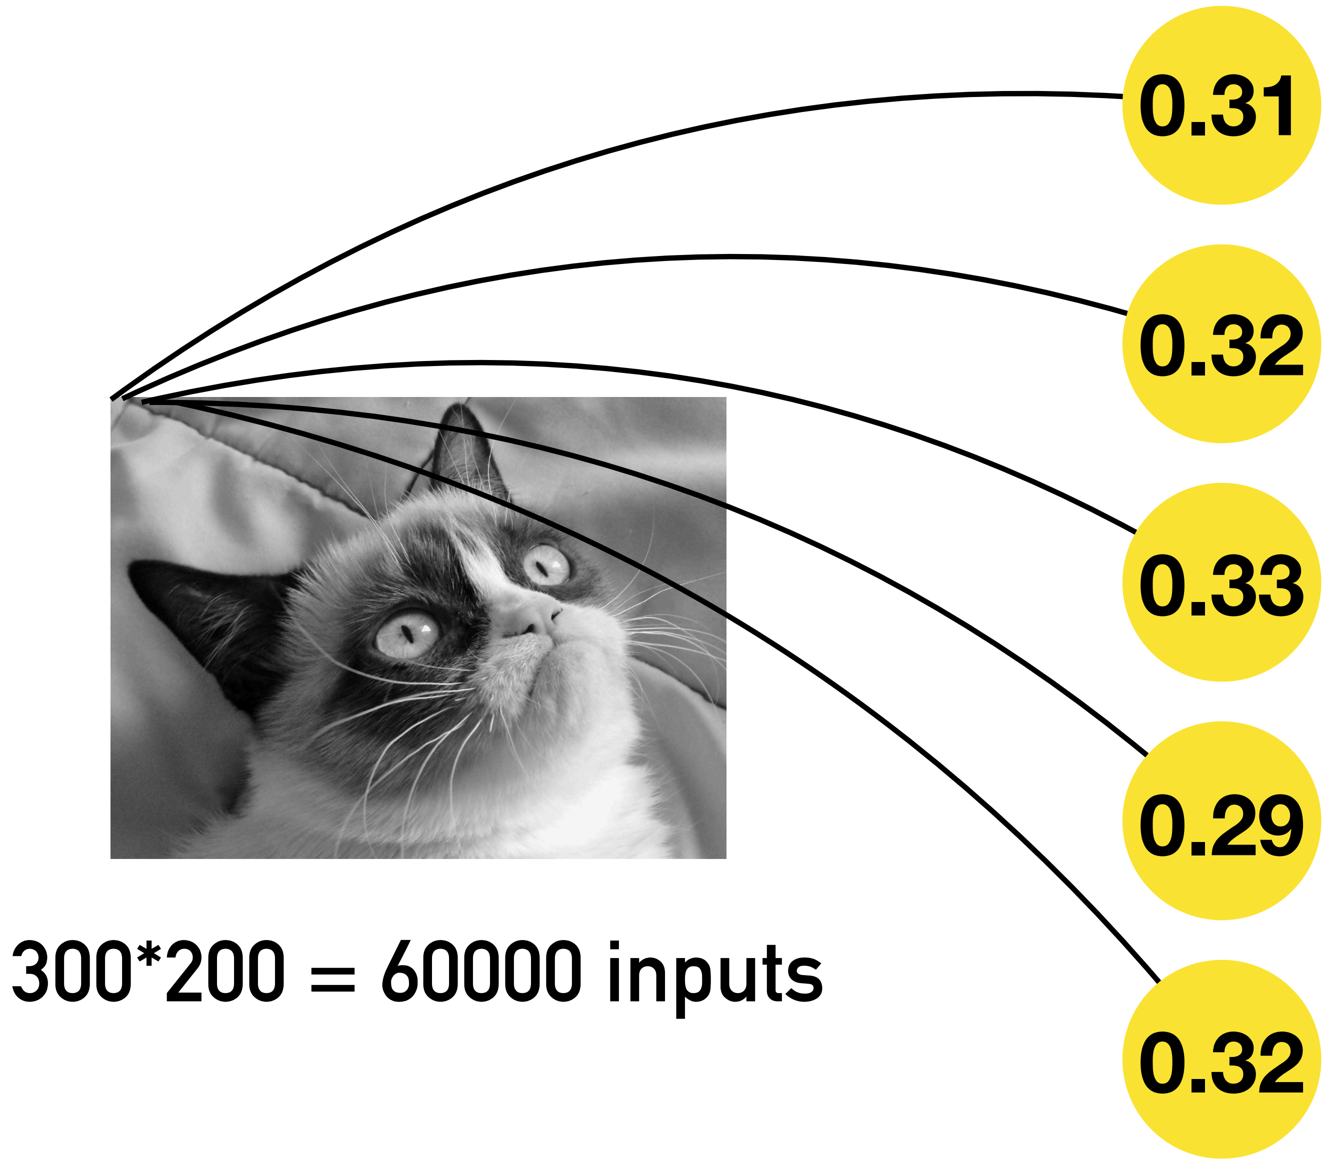

This is how a simple NN works

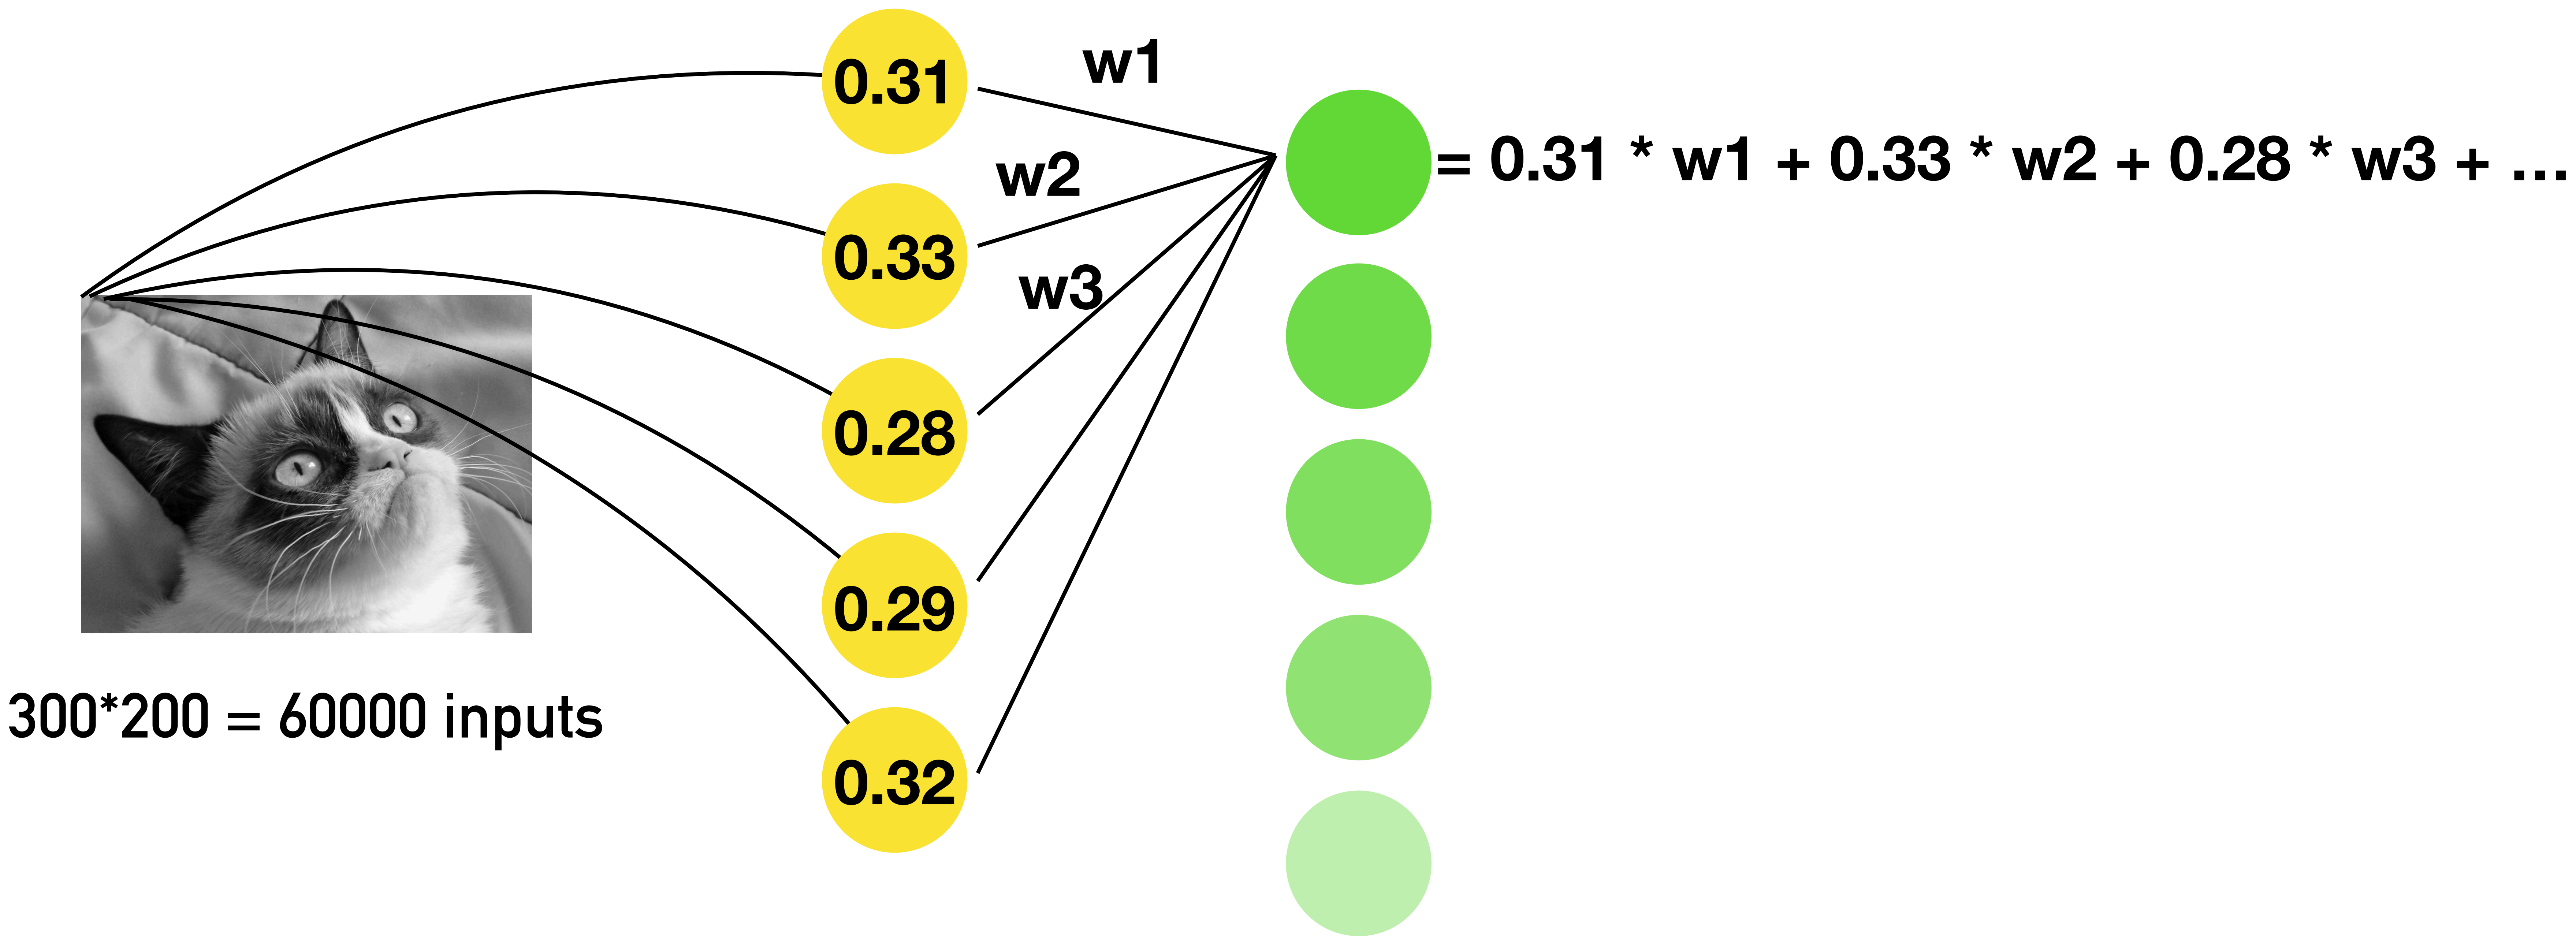

This is how a simple NN works

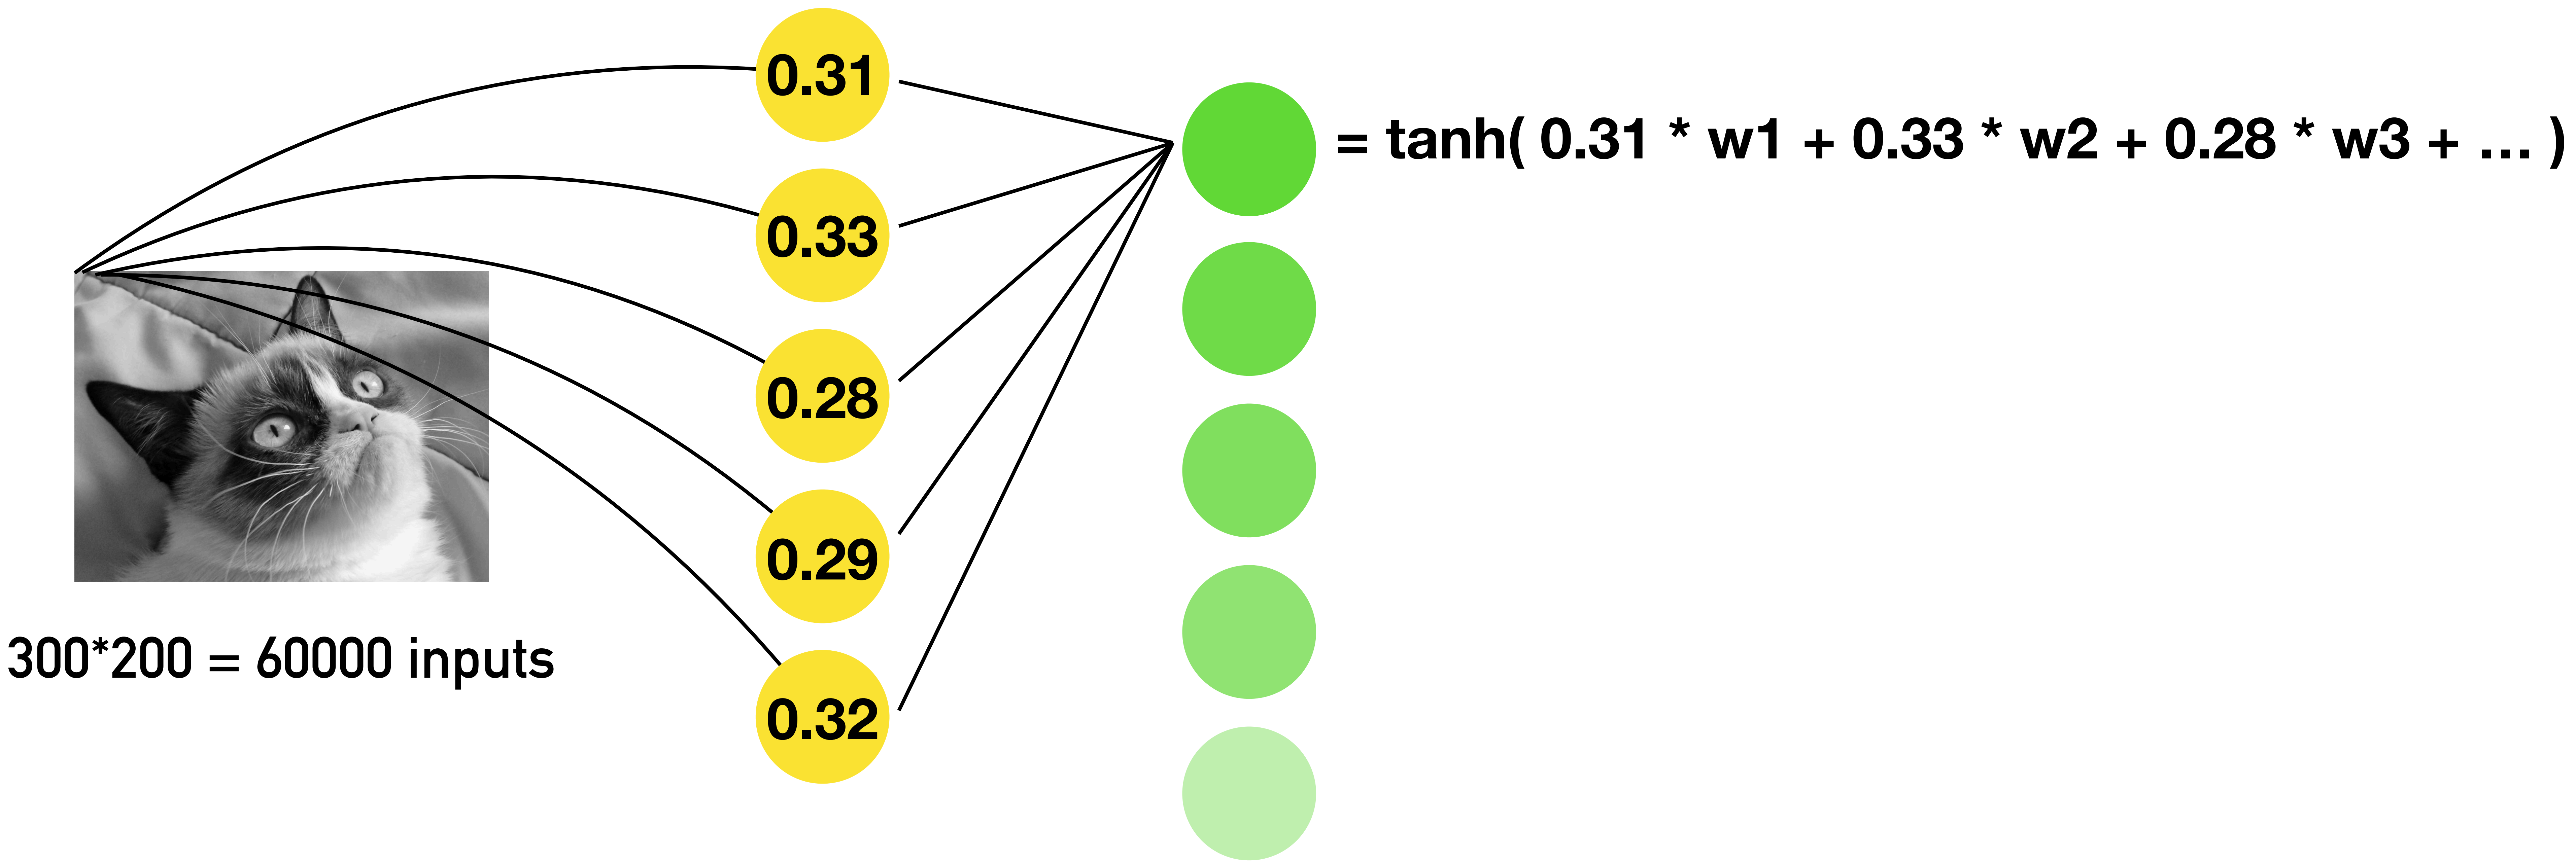

This is how a simple NN works

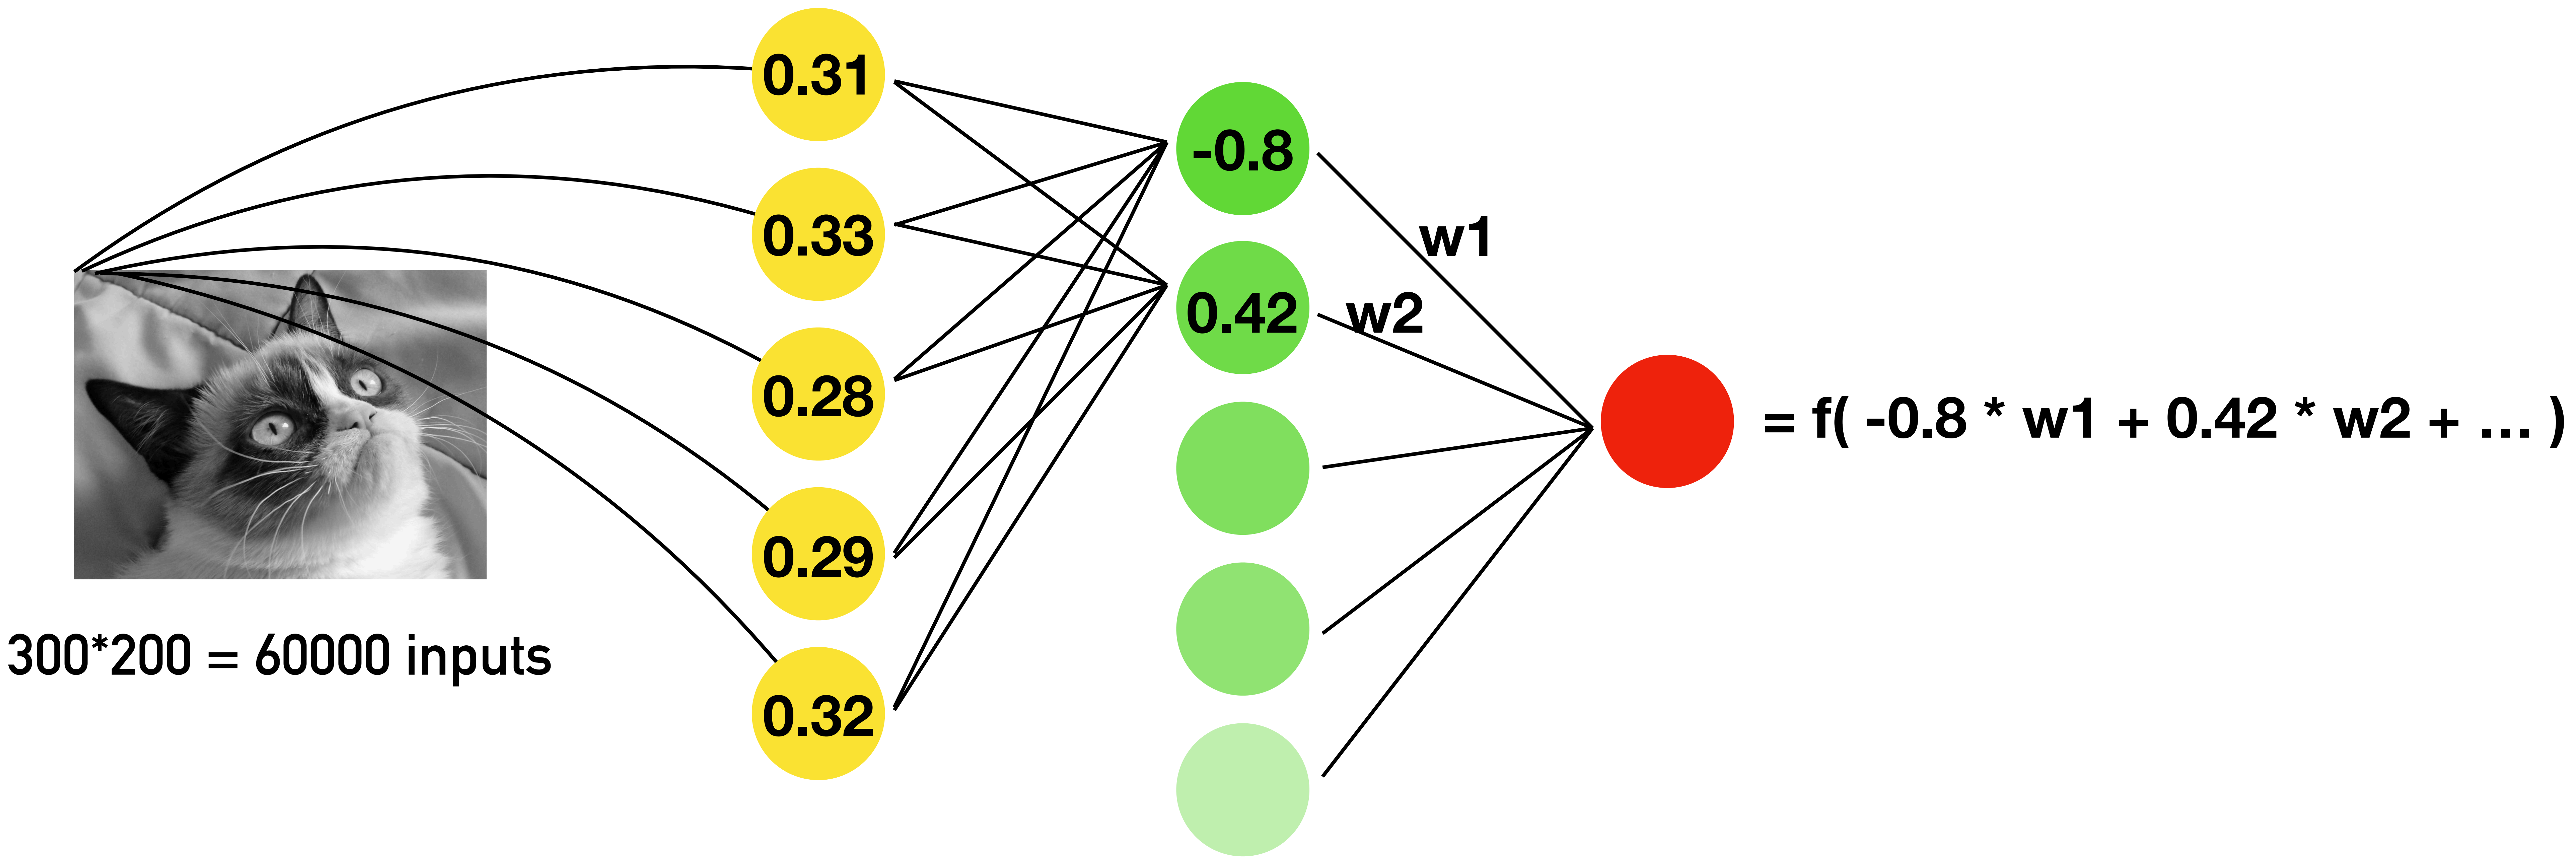

The knobs from before are the w's (the weights)

This is the entire essence of (artificial) neural networks

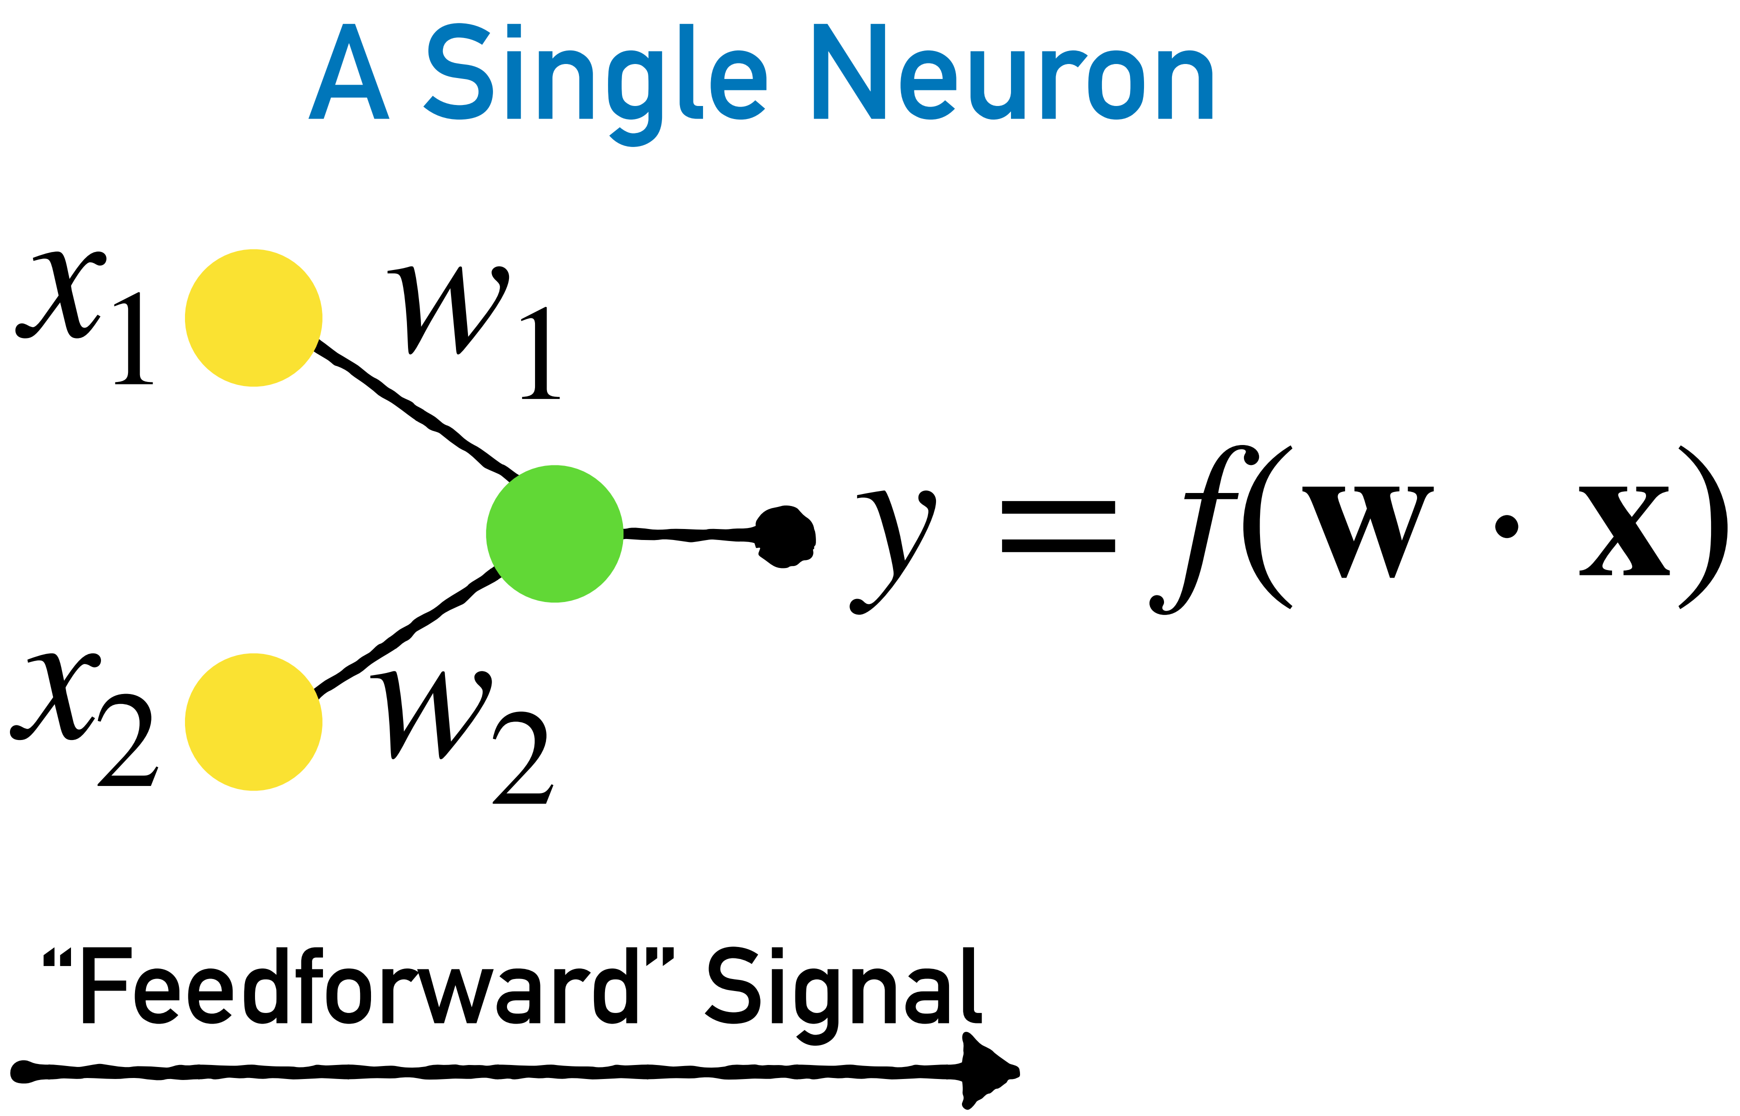

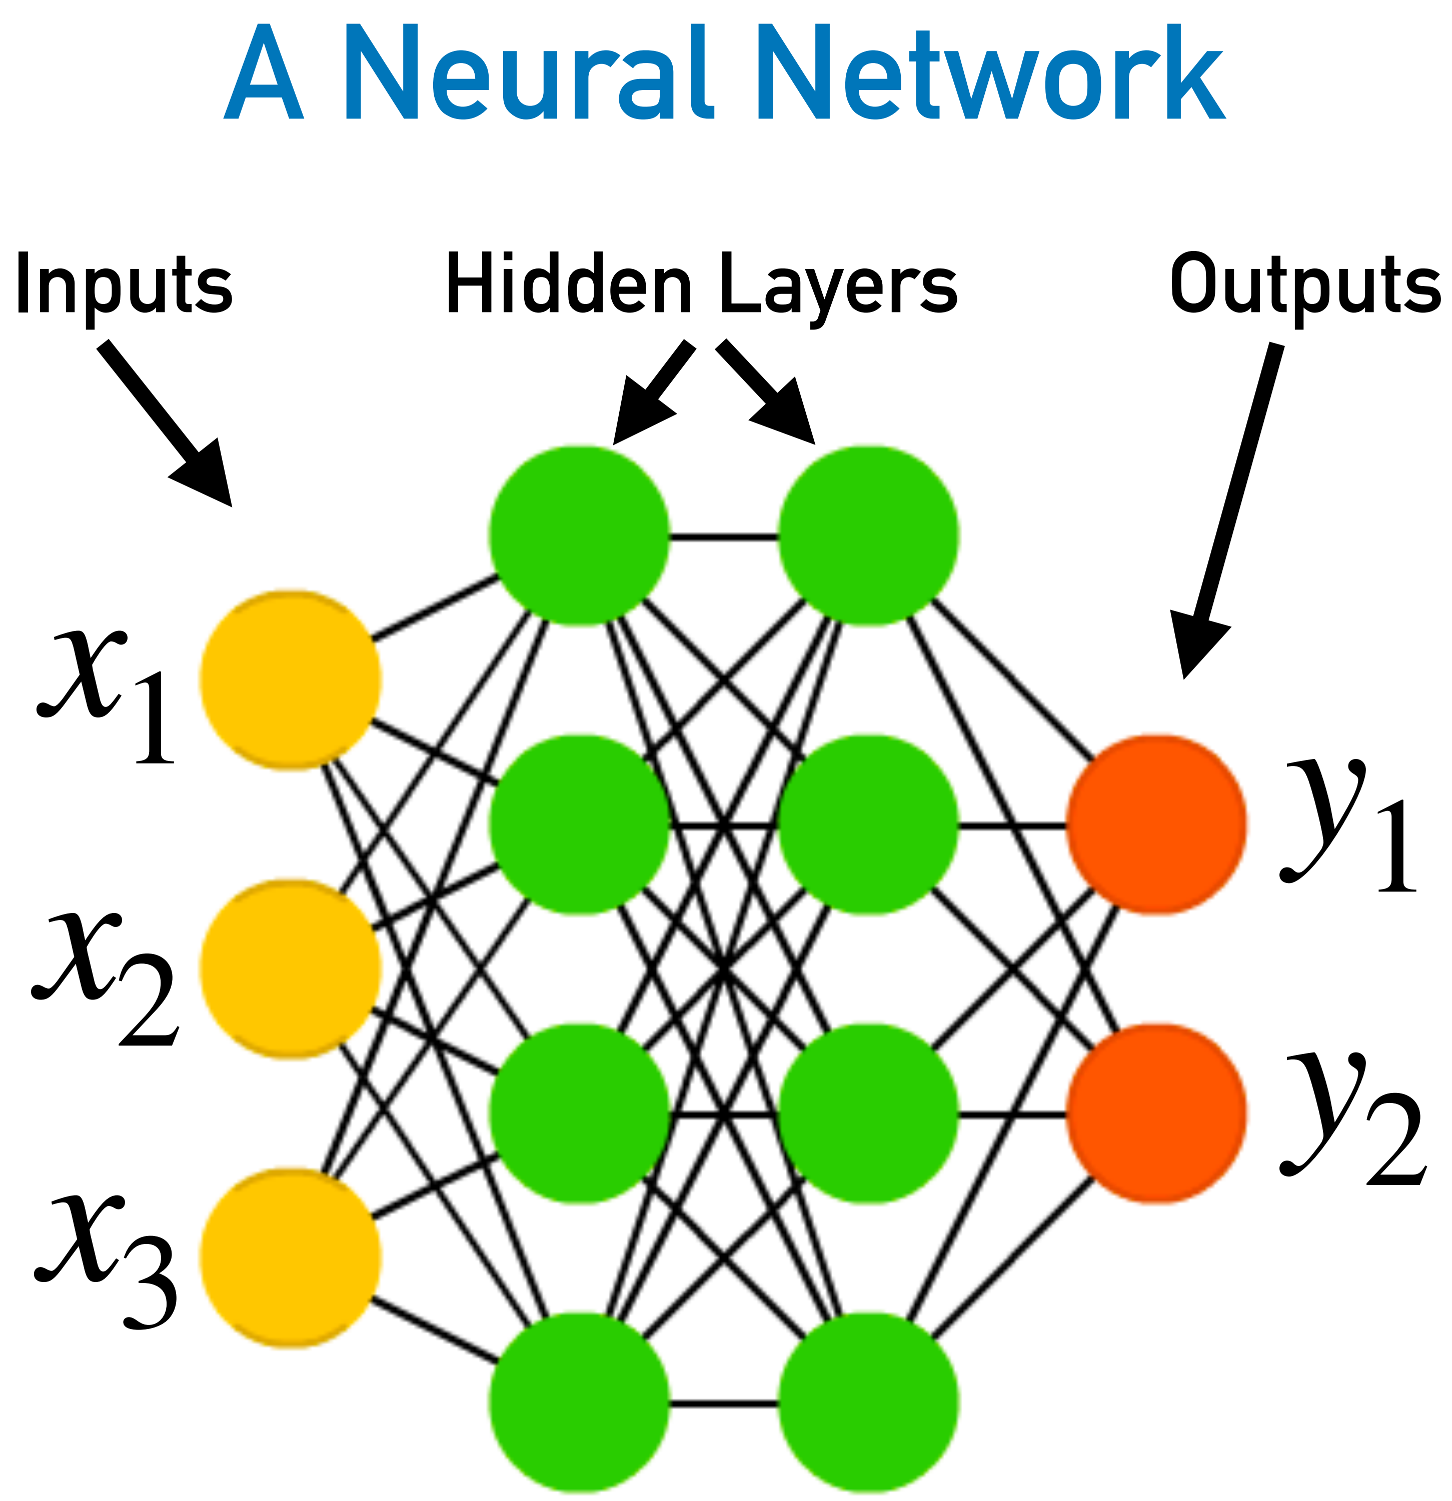

Often
$y_1 + y_2 = 1$

A Neural Network is a highly non-linear, parameterised function $y(\mathbf{x}, \{\mathbf{w}\})$

Learning = updating the weights to minimise the loss-function

Input	Answer/Target	Network Prediction	Network Prediction
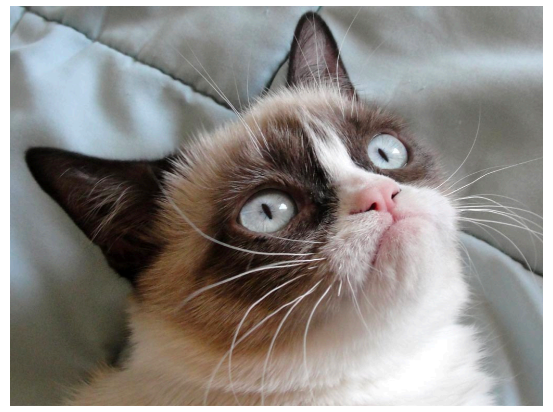	1	0	1
	0	1	0
	1	0	1
	0	1	0
		Error: 4	Error: 0

$$\mathscr{L} \sim \sum (y_{true} - y_{predicted})^2$$

Can be changed with the weights/knobs!

A workflow for neural networks looks like this:

Step 1 Given: A dataset with many inputs $\mathbf{x}$ and corresponding outputs $y_{\text{true}}(\mathbf{x})$

Step 2

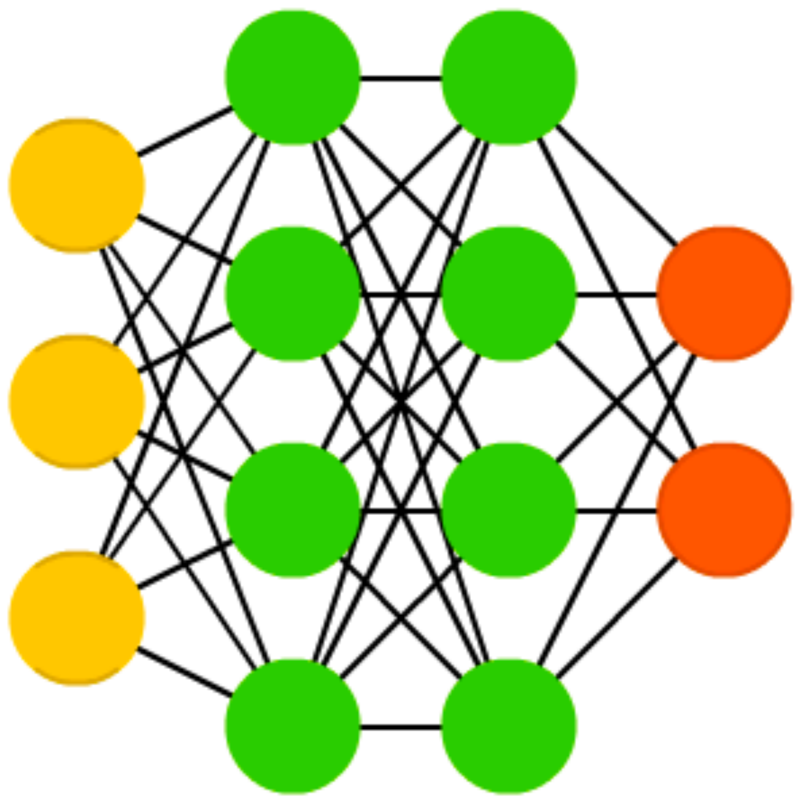

$$\mathscr{L}(\mathbf{w}, \{\mathbf{x}\}) = \sum_{\mathbf{x} \in \{\mathbf{x}\}} \left(y_{\text{true}}(\mathbf{x}) - y_{\text{network}}(\mathbf{w}, \mathbf{x}) \right)^2$$

Gradient Descent / Backpropagation $\mathbf{w} \rightarrow \mathbf{w} - \nabla_{\mathbf{w}} \mathscr{L}$

Step 3 Generalize. Use the network to infer (predict) the right output for new inputs

Neural Networks come in many topologies

https://www.asimovinstitute.org/neural-network-zoo/

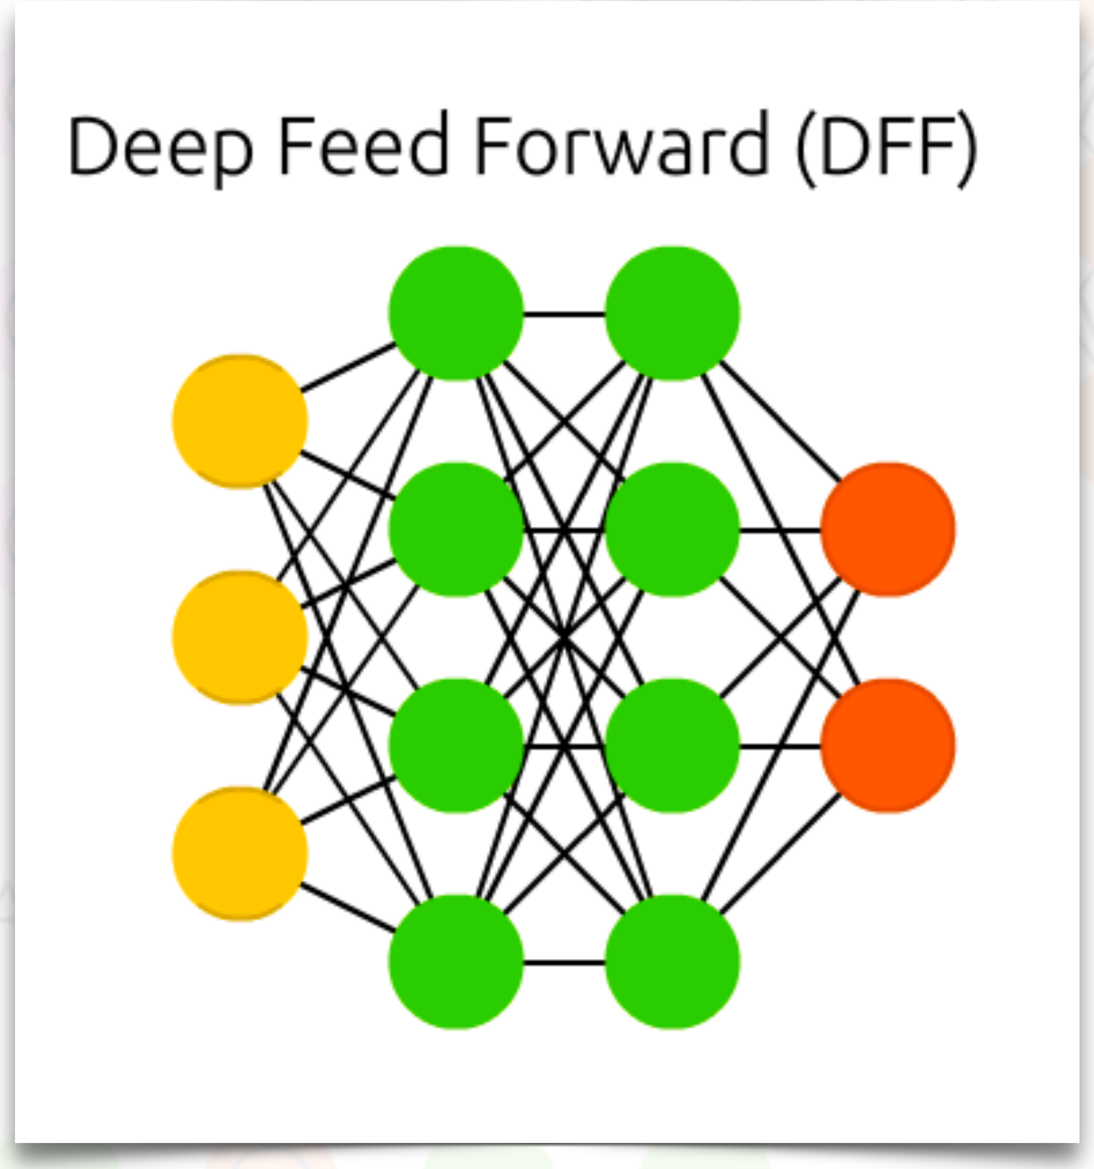

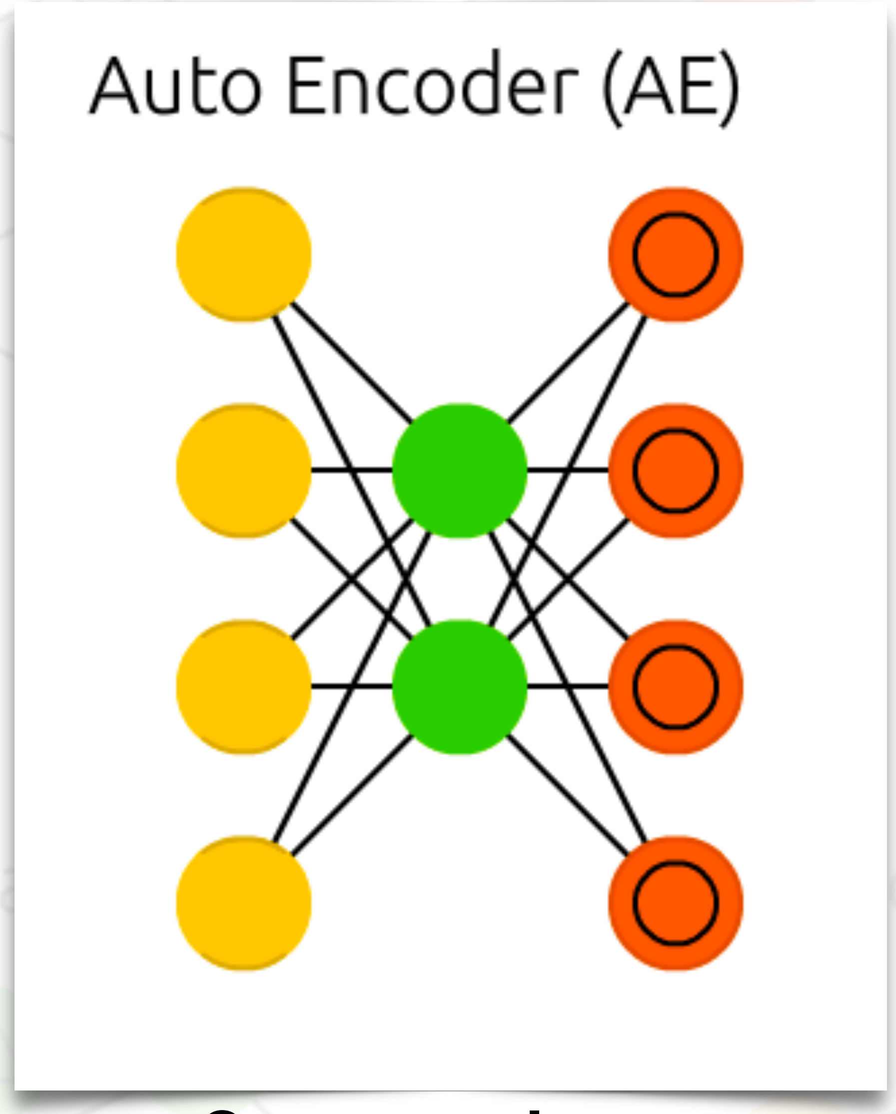

Output = Input

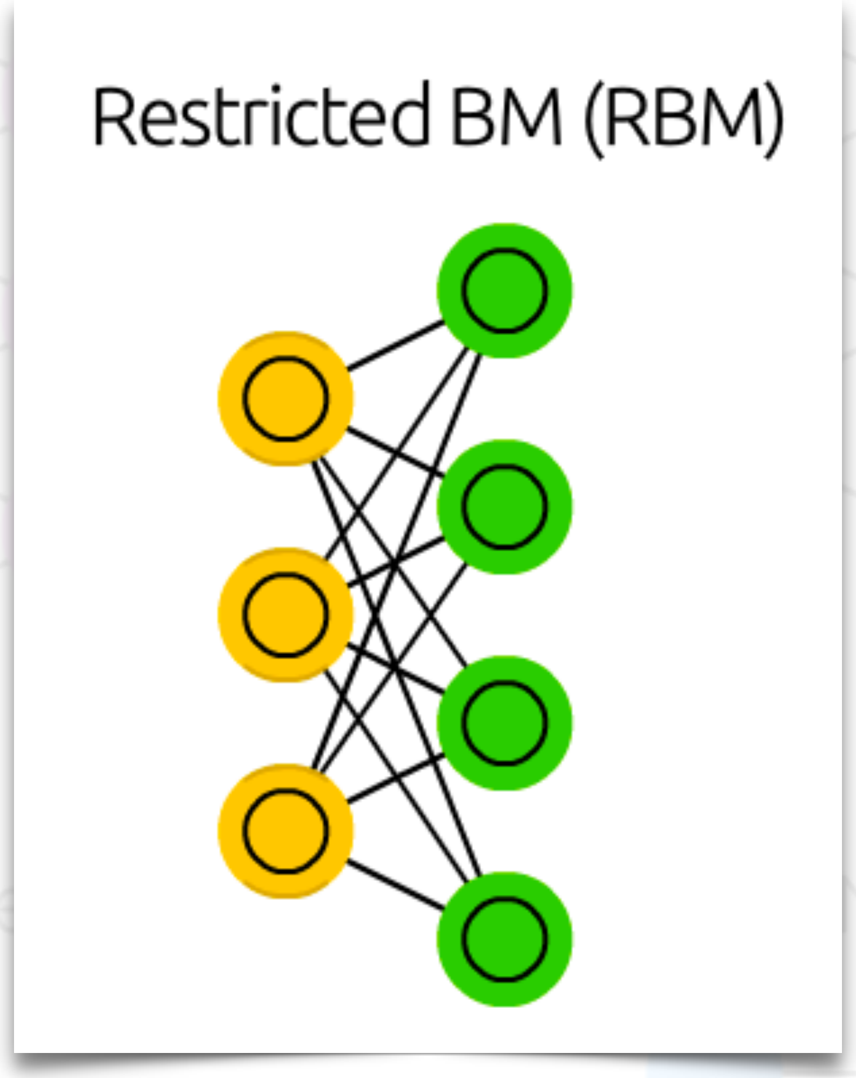

Good at learning $p(\mathbf{x})$

"Distance" between prob. distributions

Van Veen, F. & Leijnen, S. (2019). The Neural Network Zoo. Retrieved from https://www.asimovinstitute.org/neural-network-zoo

A quick summary

A Neural Network is a machine with many parameters (knobs), that we can tune (train) so that it reproduces the answers we want

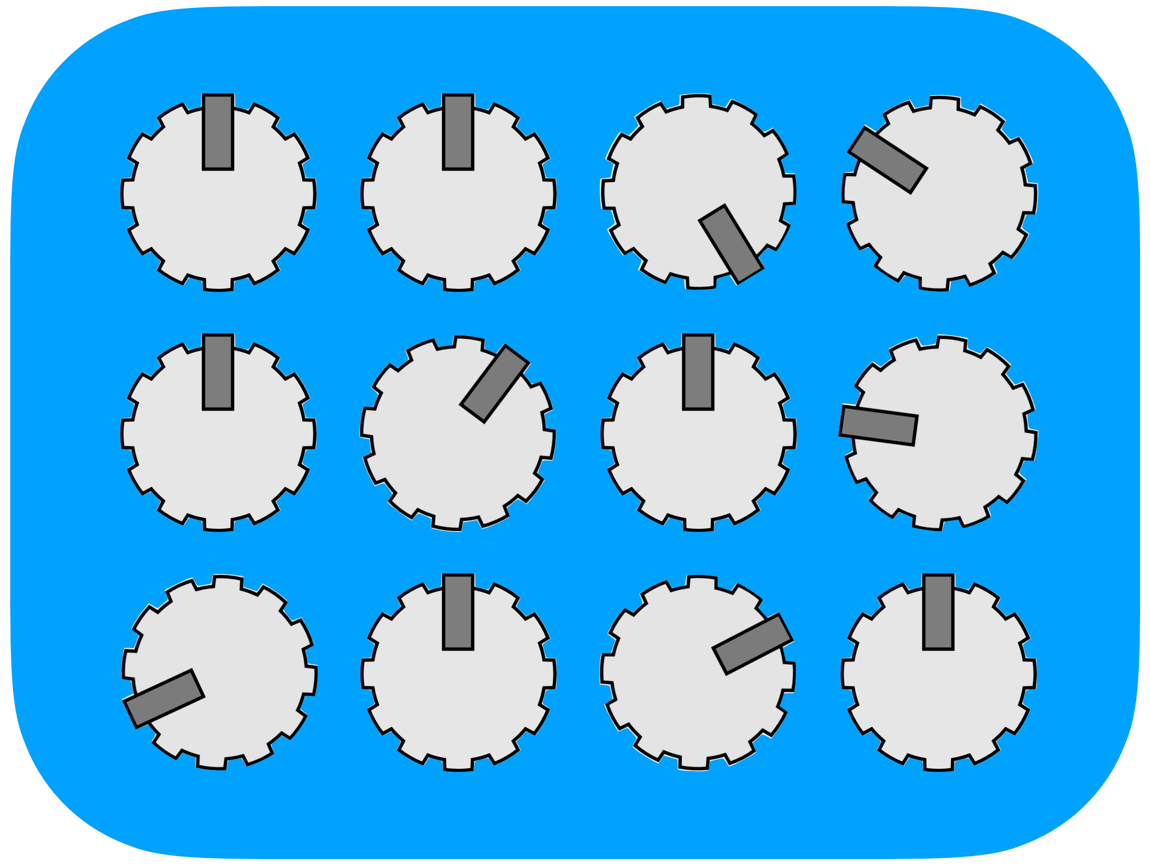

Given enough parameters, we can fit any function we want

(For example, we can fit the 'is_this_a_cat?' function)
(Or, we can fit a 'turn_random_noise_into_a_face' function)
(... etc)

There are roughly three types of ML

Supervised Learning
Learning from examples

$\{x_1, y_1\} = \{$ 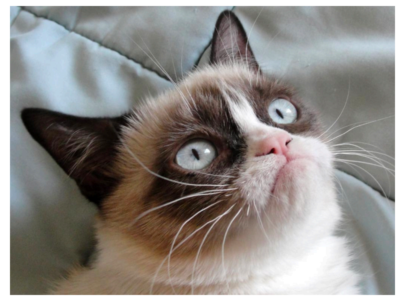 $, 1\}$

$\{x_2, y_2\} = \{$  $, 0\}$

Learn $p(y \mid x)$

Classification

Unsupervised Learning
Learning about examples

$x_1 =$ 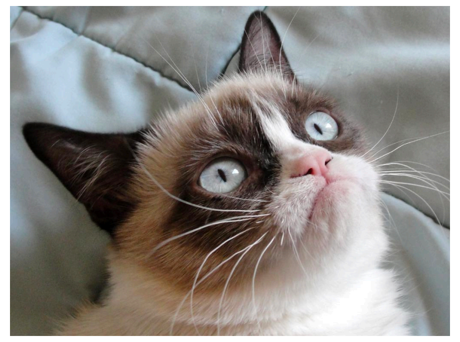

$x_2 =$ 

Learn $p(x)$

(Draw samples to generate!)

Reinforcement Learning
Learning from feedback

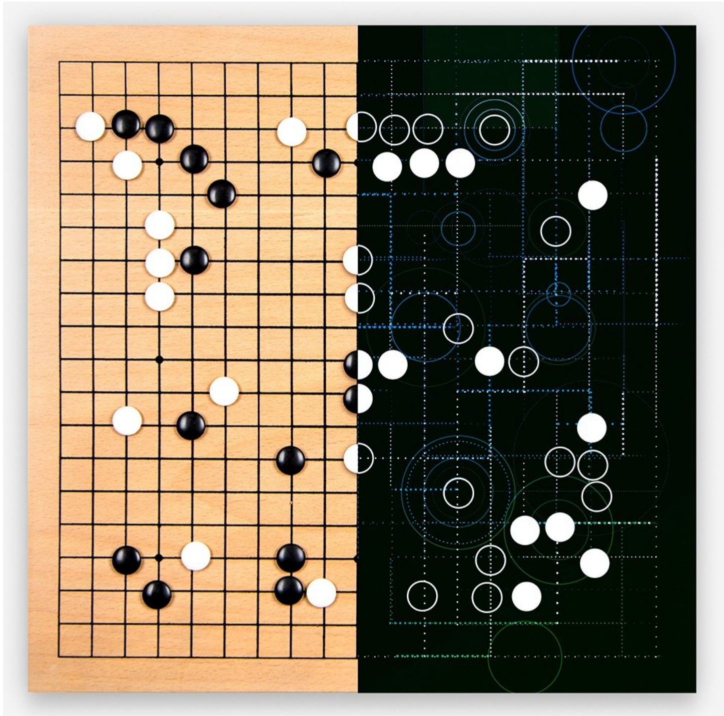

Learn a policy,
(best action in a given state s)

Learn $\pi(s)$

Sutton&Barto

Each of these types has a use in physics

Supervised Learning

Learning from examples

Picture of a galaxy -> which type?

LHC collisions -> which particles?
LHC -> interesting collision?

Material -> superconductor?

Learn $p(y|x)$

Classification

Unsupervised Learning

Learning about examples

Generate more superconductors?

Run EXPERIMENTS!

Learn $p(x)$

(Draw samples to generate!)

Reinforcement Learning

Learning from feedback

Correct errors in a quantum computer

Control EXPERIMENTS!

Learn $\pi(s)$

Sutton&Barto

Condensed Matter Physics

Studies properties of matter

How well does a piece of metal conduct?
Why are metals shiny?
How do superconductors work?
How does an insulator work?
How do we make a quantum computer?
At what temperature does a magnet stop working?
...

This is how we use ML for Quantum Physics

Supervised Learning

Finding Phase Transitions

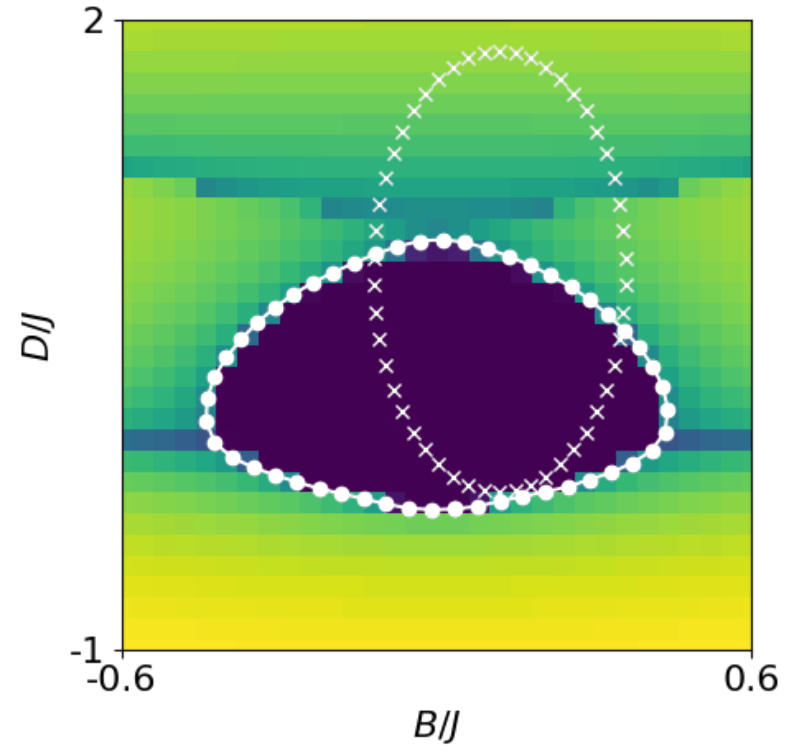

Enhancing experiments

Unsupervised Learning

Wavefunction Neural Network

Quantum State Reconstruction

$\Psi =$

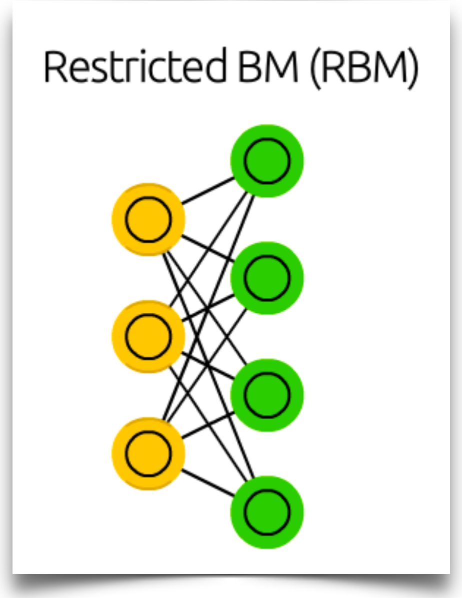

Reinforcement Learning

Correcting a quantum computer

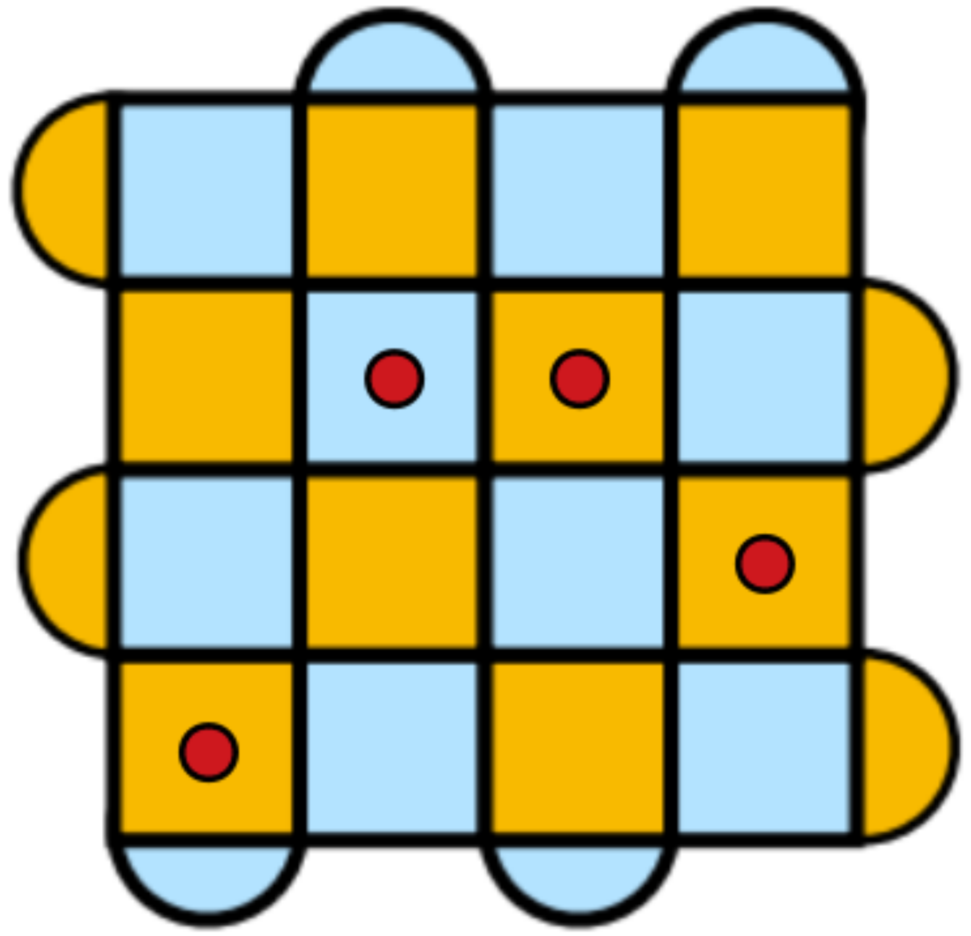

Controlling experiments

This is how we use ML for Quantum Physics

Supervised Learning

Finding Phase Transitions

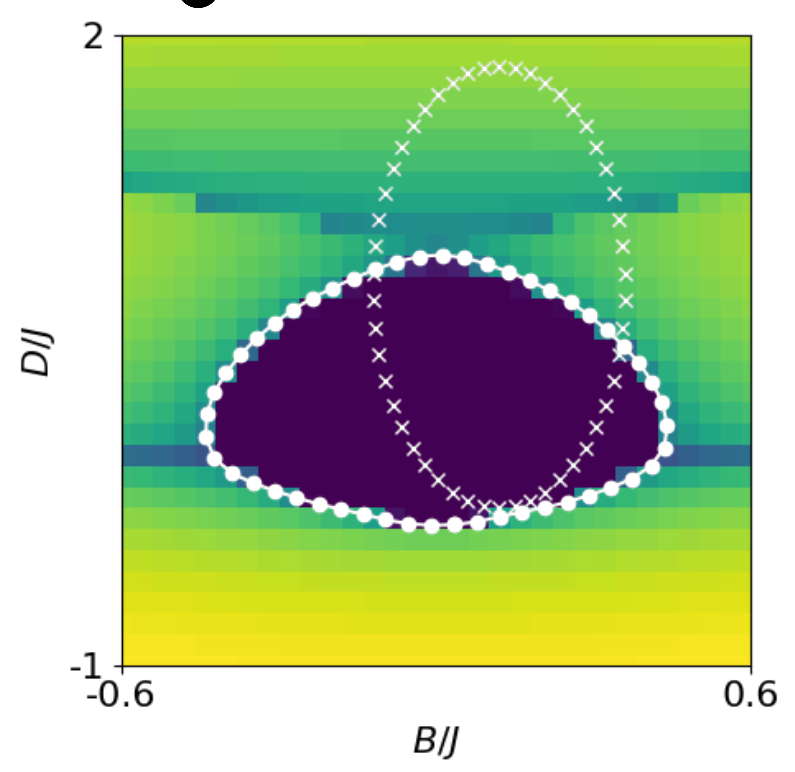

Enhancing experiments

Supervised learning can be used to find phase transitions

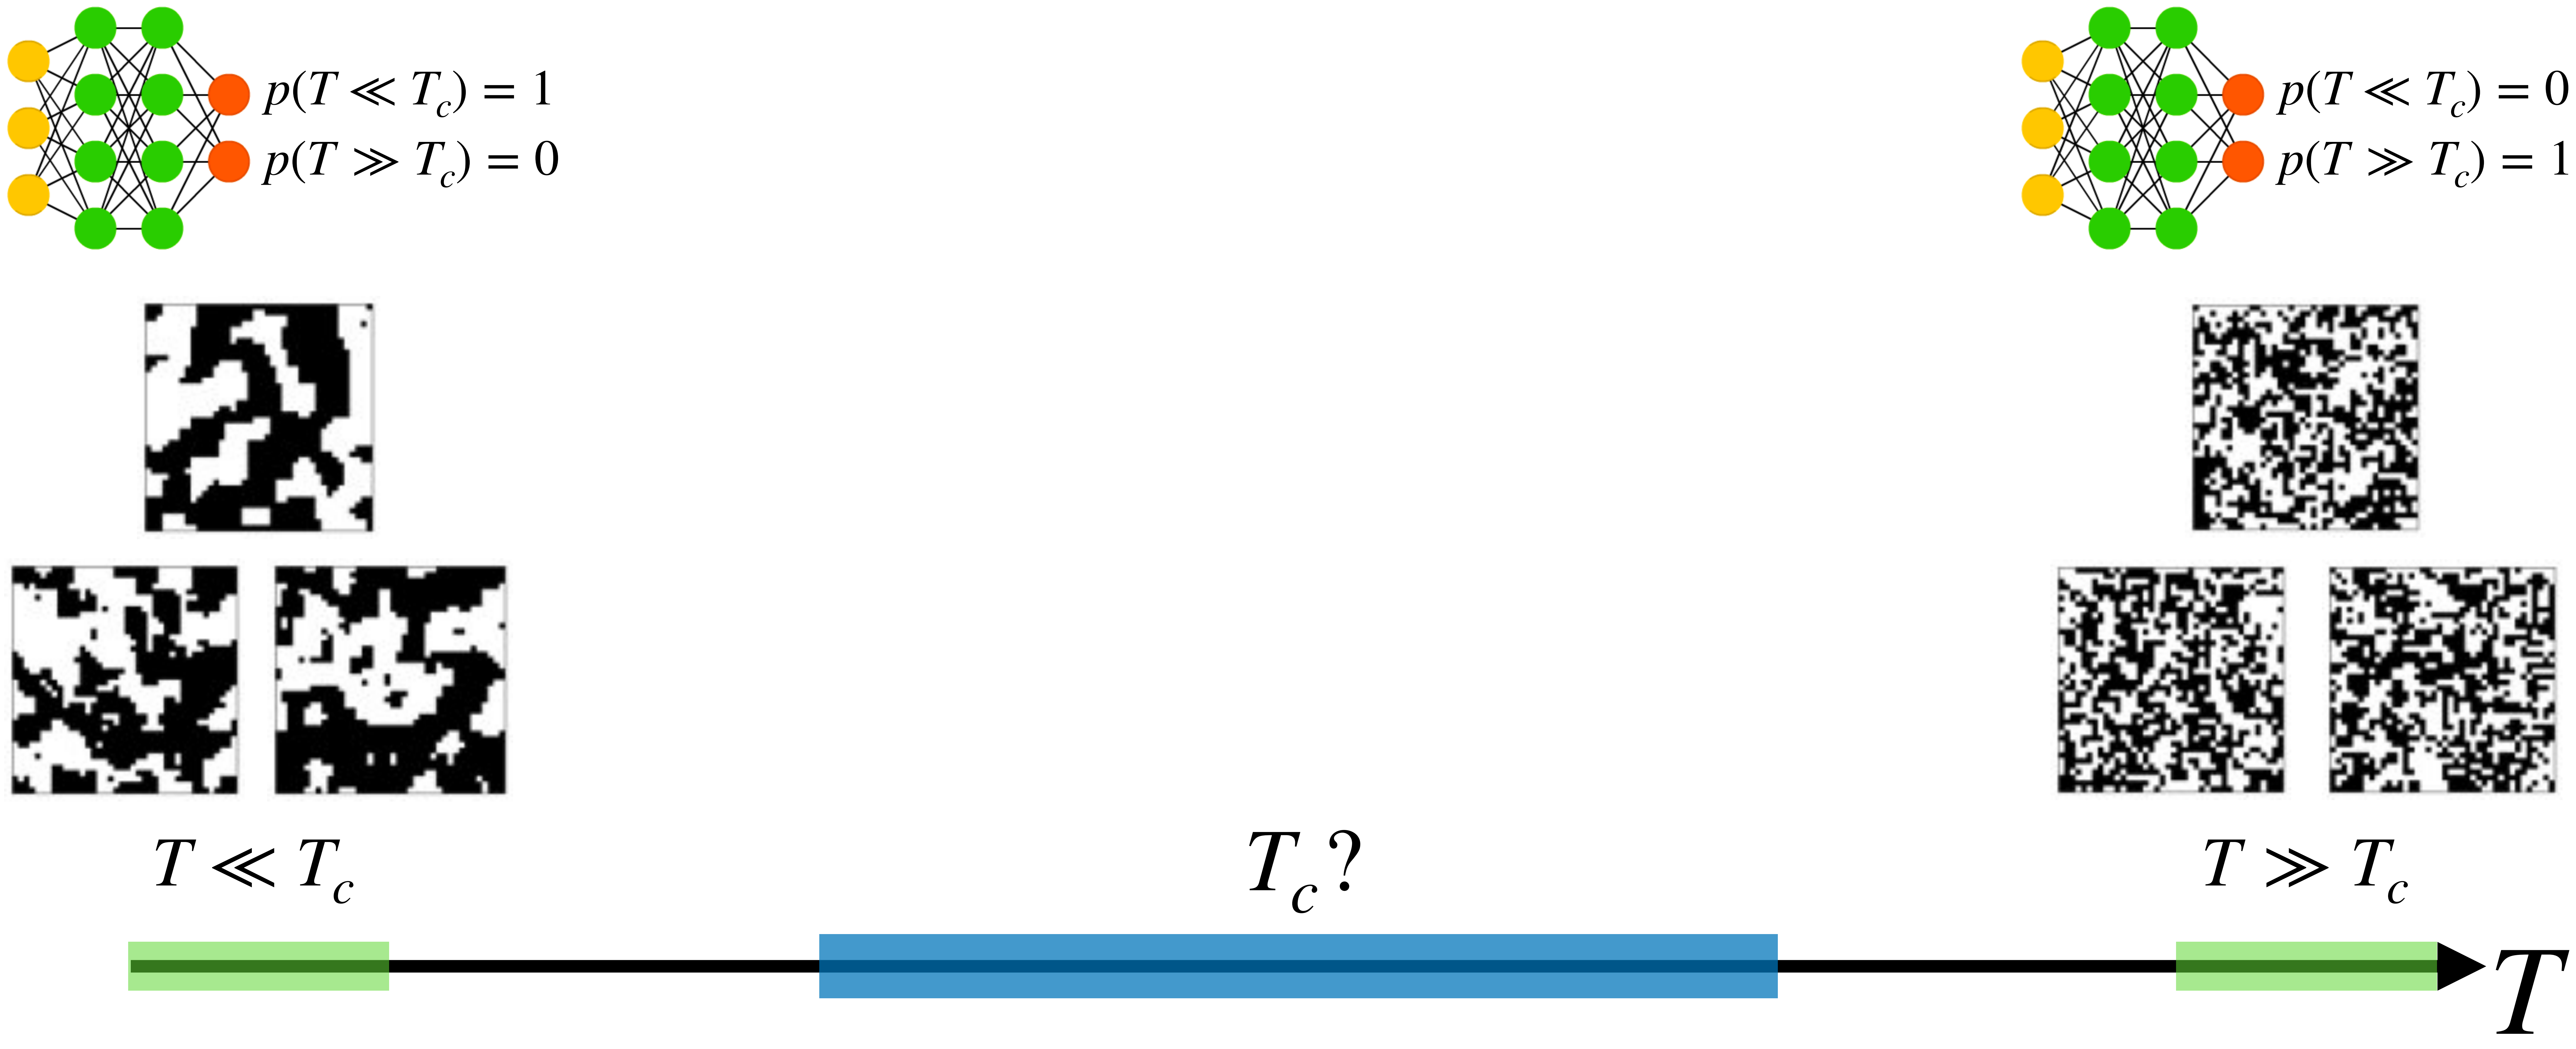

Supervised learning can be used to find phase transitions

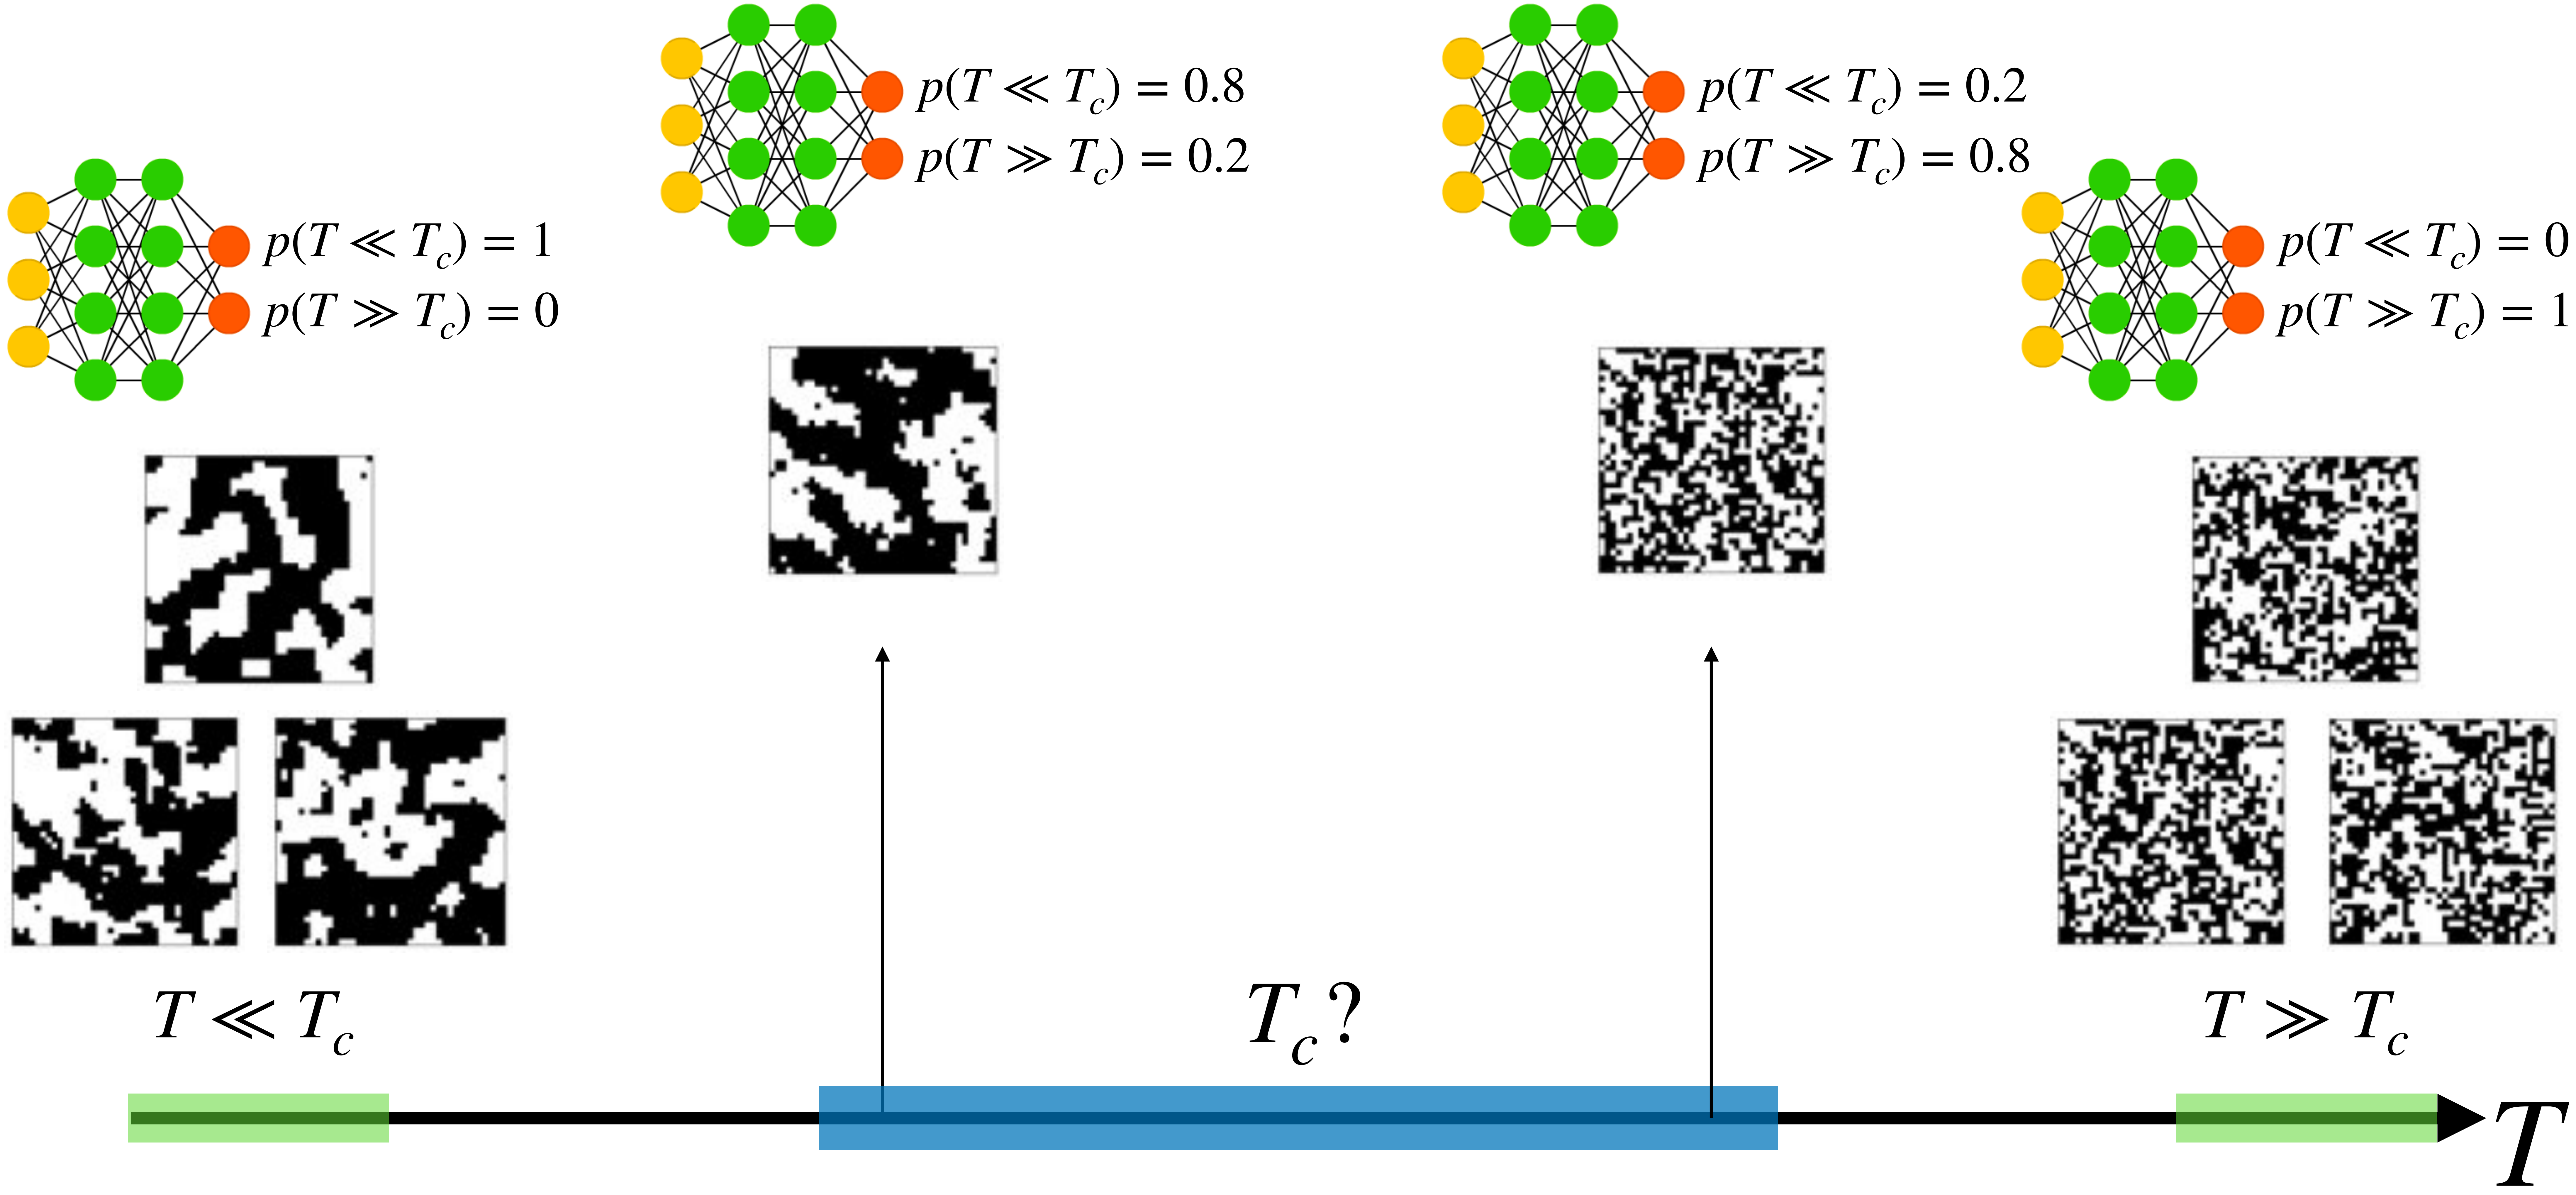

Supervised learning can be used to find phase transitions

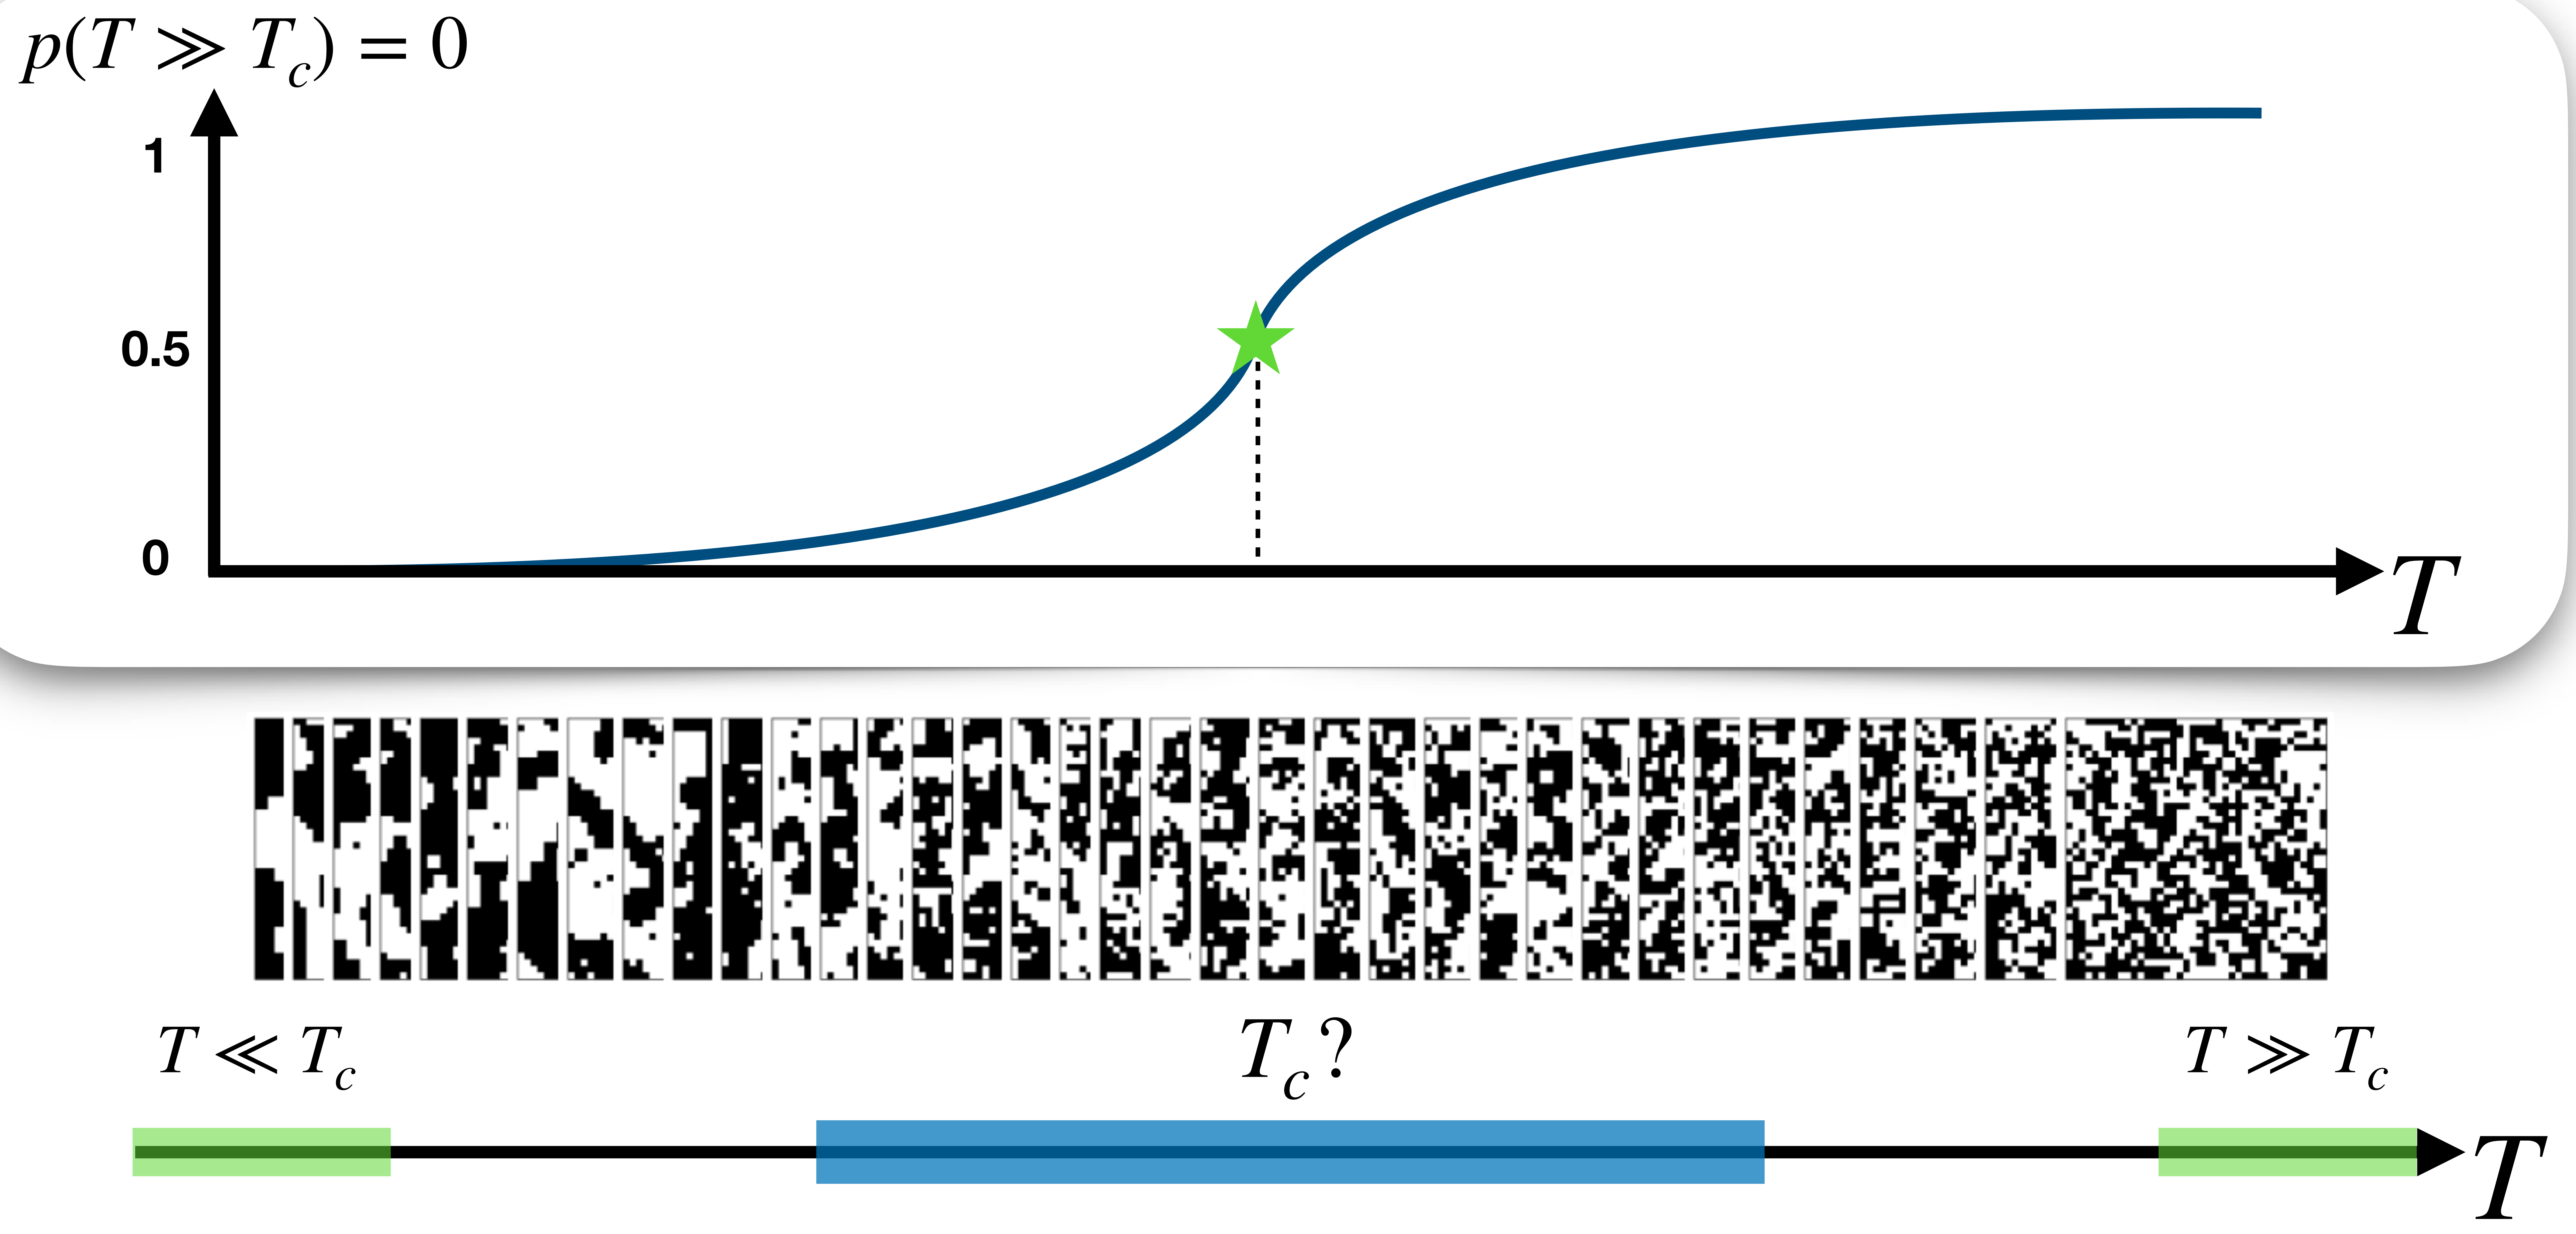

Supervised learning can make some experiments 100x faster

At QDev, Ferdinand Kuemmeth works on qubits

Quantum bit: can be part 0 and part 1 simultaneously

Reading out the qubit state: electric signals

The signal is then demodulated using techniques that are rooted in old radio-technology

Takes 10-100 microseconds!

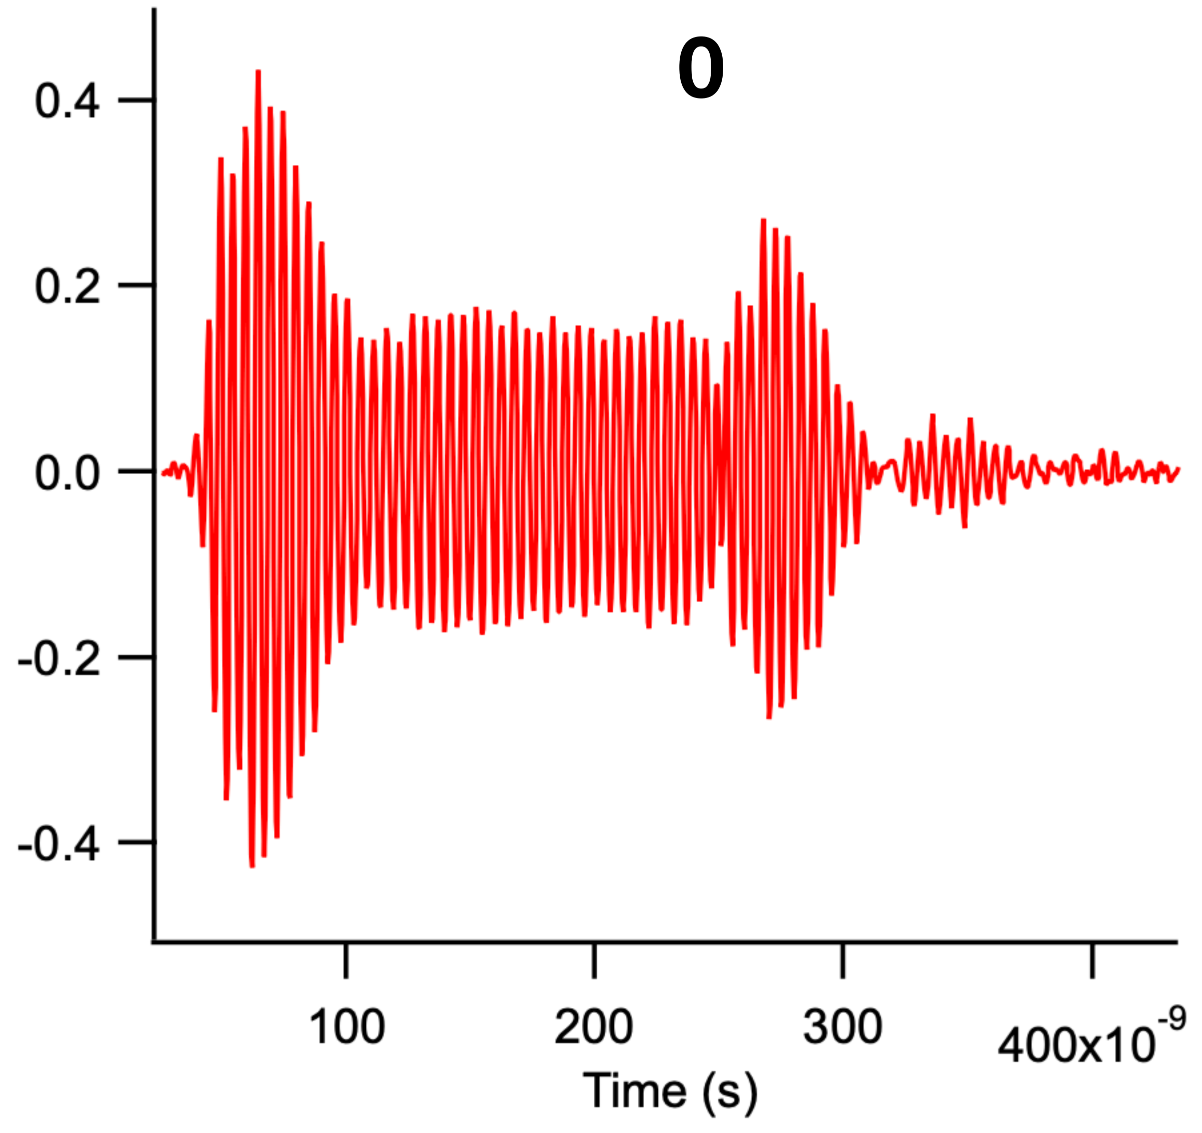

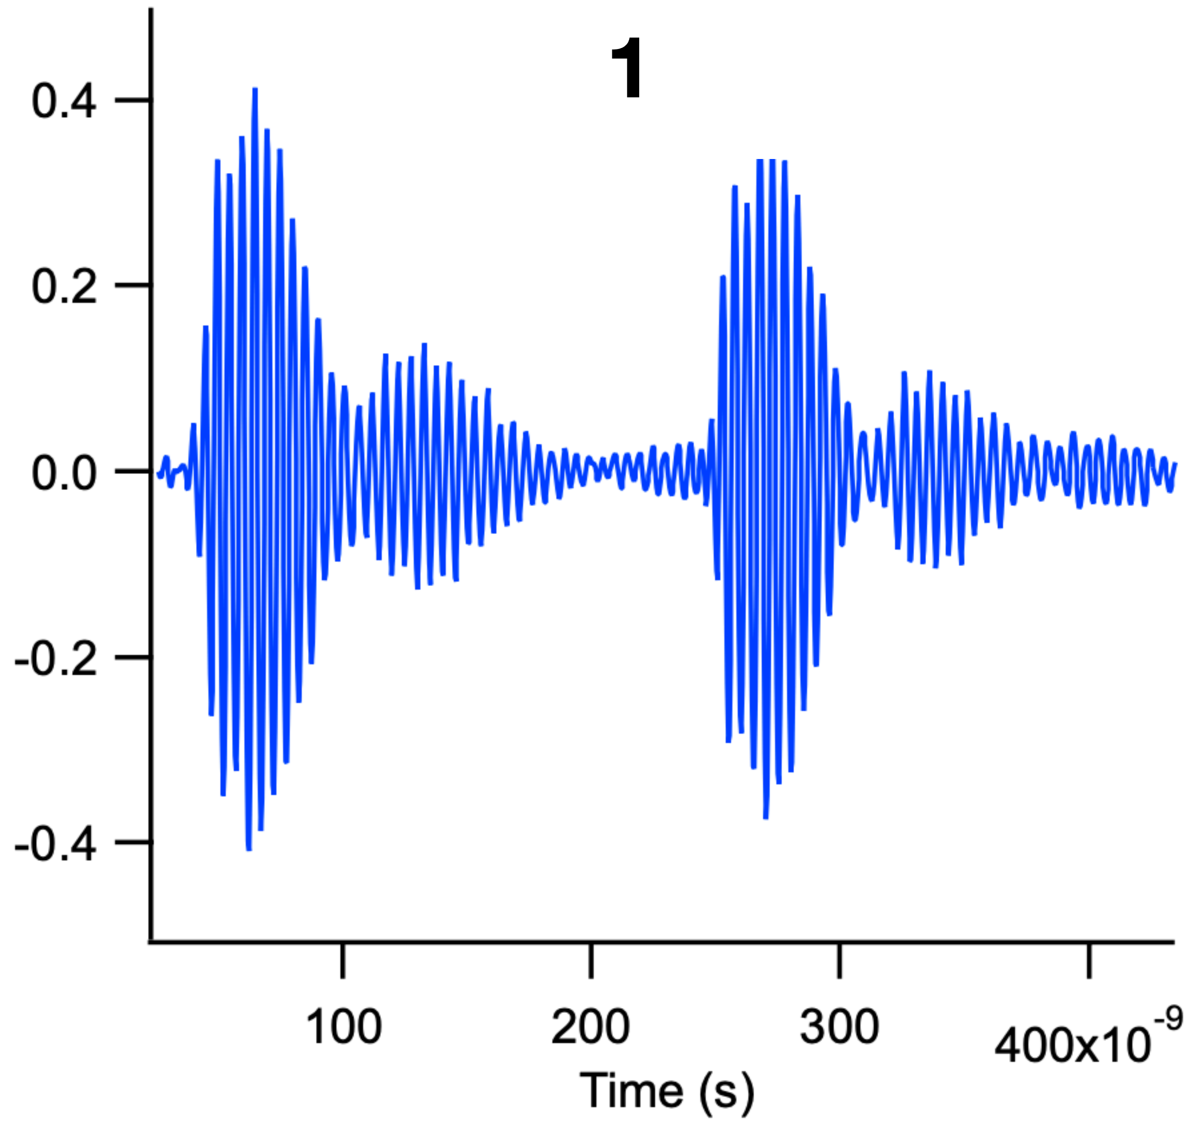

Raw data -> ~100 nanoseconds!

This is how we use ML for Quantum Physics

Supervised Learning

Finding Phase Transitions

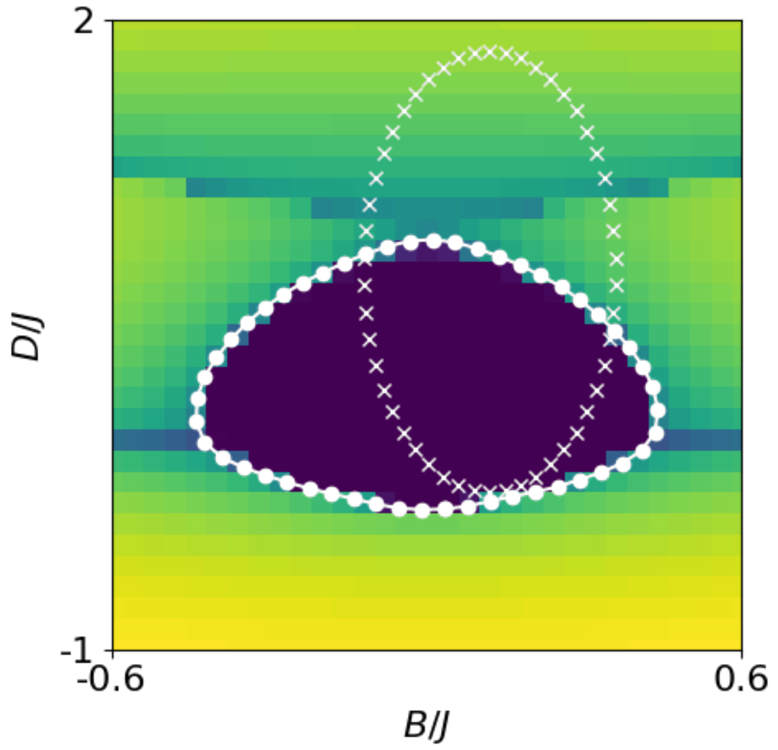

Enhancing experiments

Unsupervised Learning

Wavefunction Neural Network

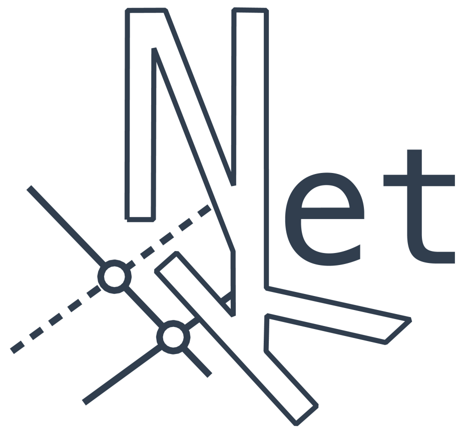

www.netket.org

Quantum State Reconstruction

$$\Psi =$$

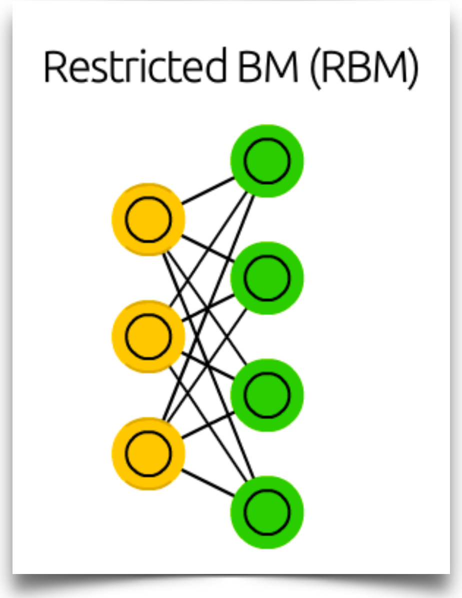

Unsupervised learning can do Quantum State Reconstruction

An experiment with 4 two-level atoms (each can be in quantum state 0 or 1)

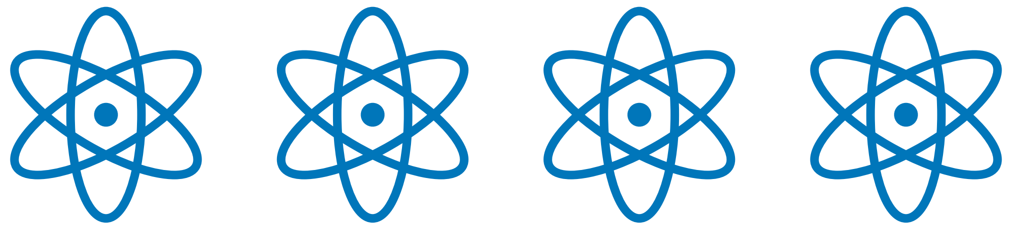

The full system can be a superposition of all 16 possible states

0 0 0 0
0 0 0 1
0 0 1 0
0 0 1 1
. . .

The full state of this system is called the wavefunction

$$\Psi(a_1, a_2, a_3, a_4)$$

Unsupervised learning can do Quantum State Reconstruction

An experiment with 4 two-level atoms (each can be in quantum state 0 or 1)

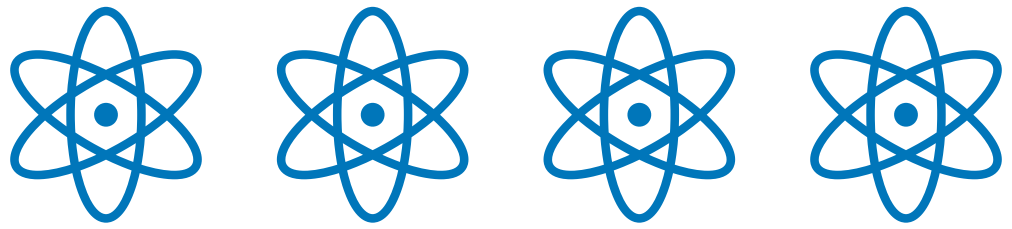

Perform the experiment many times, and record which configuration we get

0 1 0 1
1 1 0 1
1 0 0 1
0 1 0 1
0 1 0 1

Projective measurements

Every time we look, we force the atoms to choose; Schrödingers cat

Question: can we learn $\Psi(a_1, a_2, a_3, a_4)$?

Unsupervised learning can do Quantum State Reconstruction

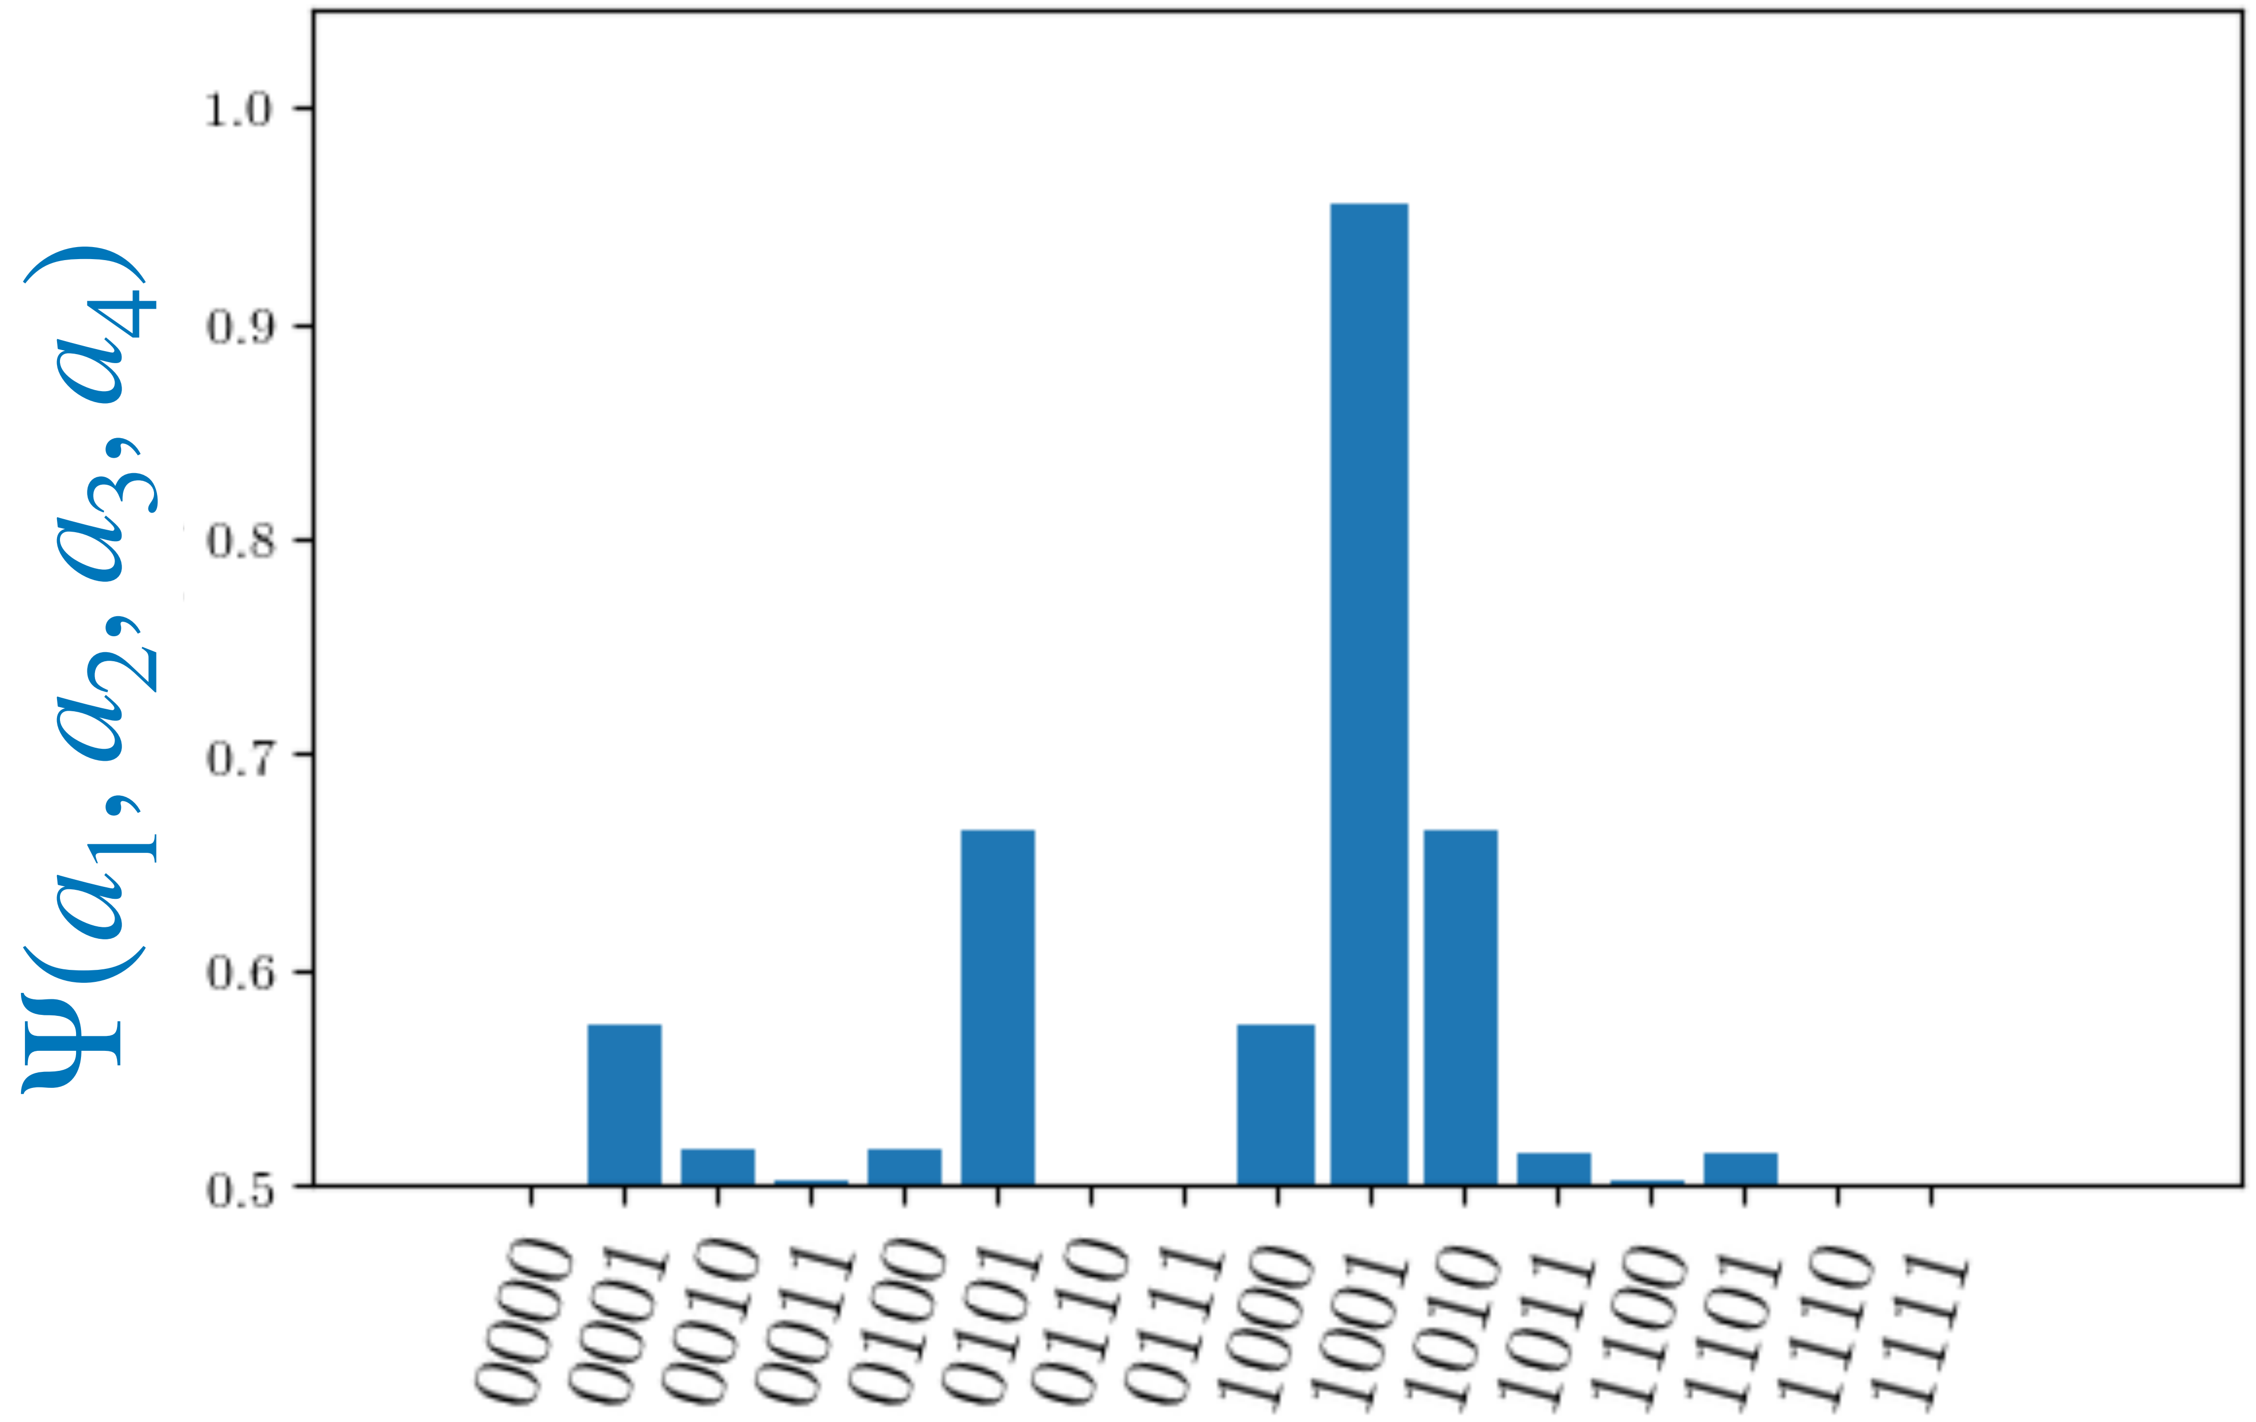

Histogram does not scale!

Exponential number of measurements!

2^4 = 16 2^64 > grains of sand on earth

$\Psi =$

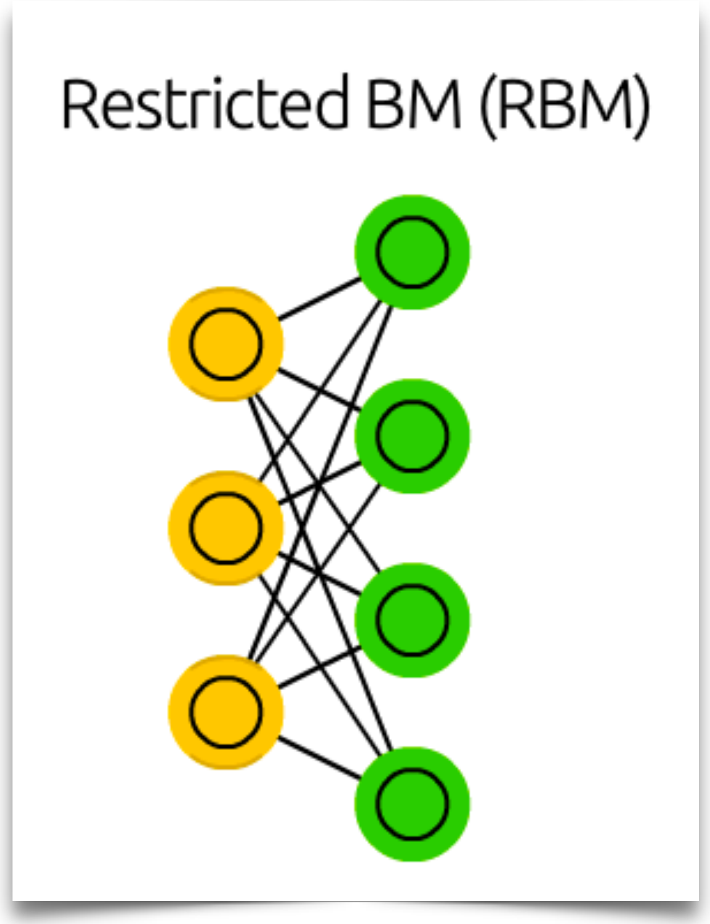

Trade a complex (impossible) measurement for many simple ones

This is how we use ML for Quantum Physics

Supervised Learning

Finding Phase Transitions

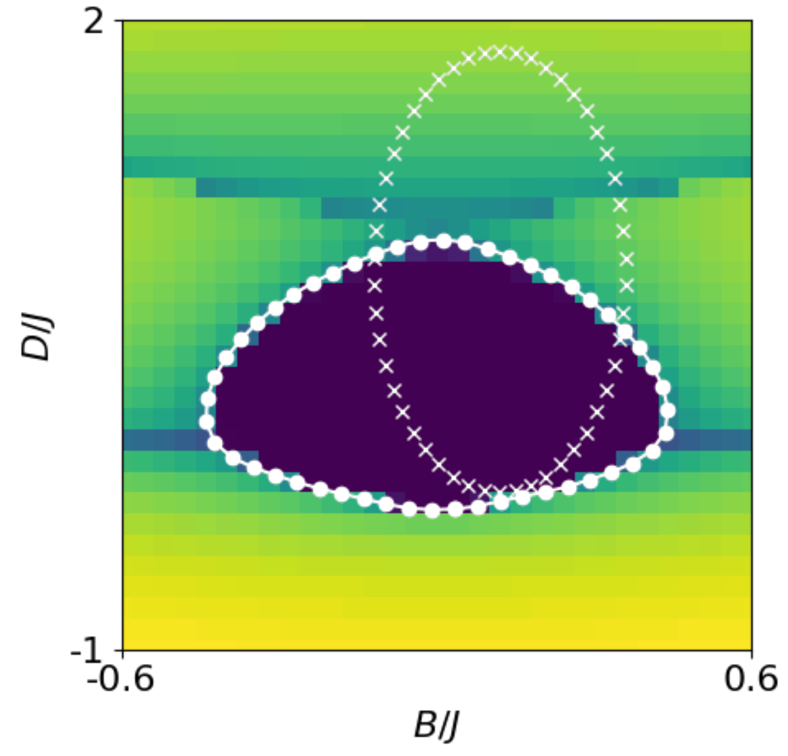

Enhancing experiments

Unsupervised Learning

Wavefunction Neural Network

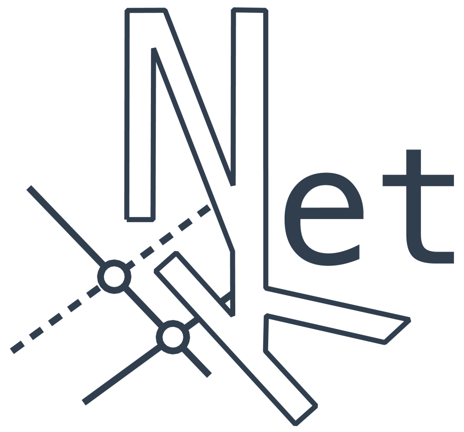

www.netket.org

Quantum State Reconstruction

$\Psi =$ 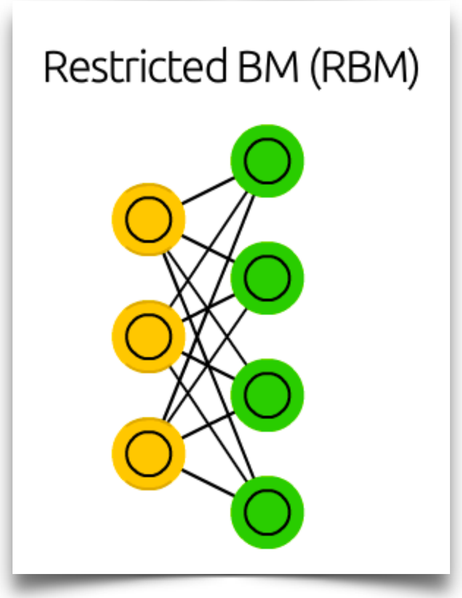

Reinforcement Learning

Correcting a quantum computer

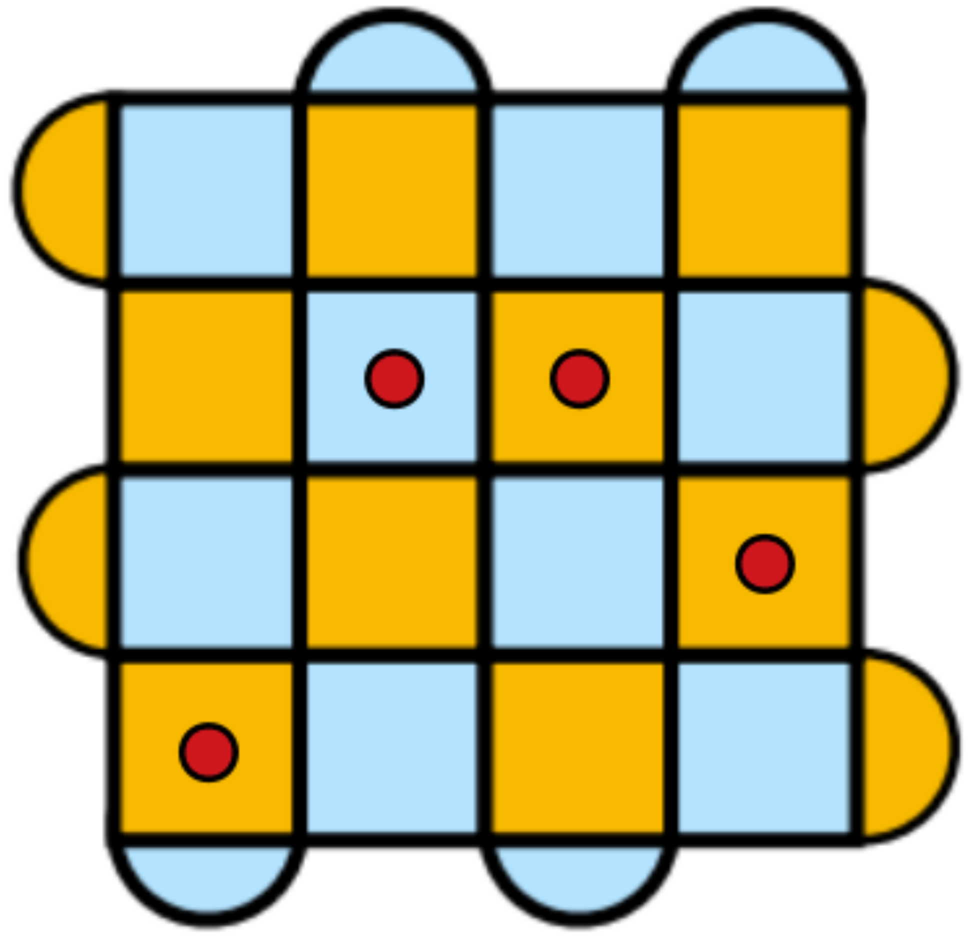

Controlling experiments

Reinforcement learning is about strategies (policies)

MuZero (the successor to AlphaZero and AlphaGo)

Image Credit: DeepMind

A simple quantum computer with one logical qubit

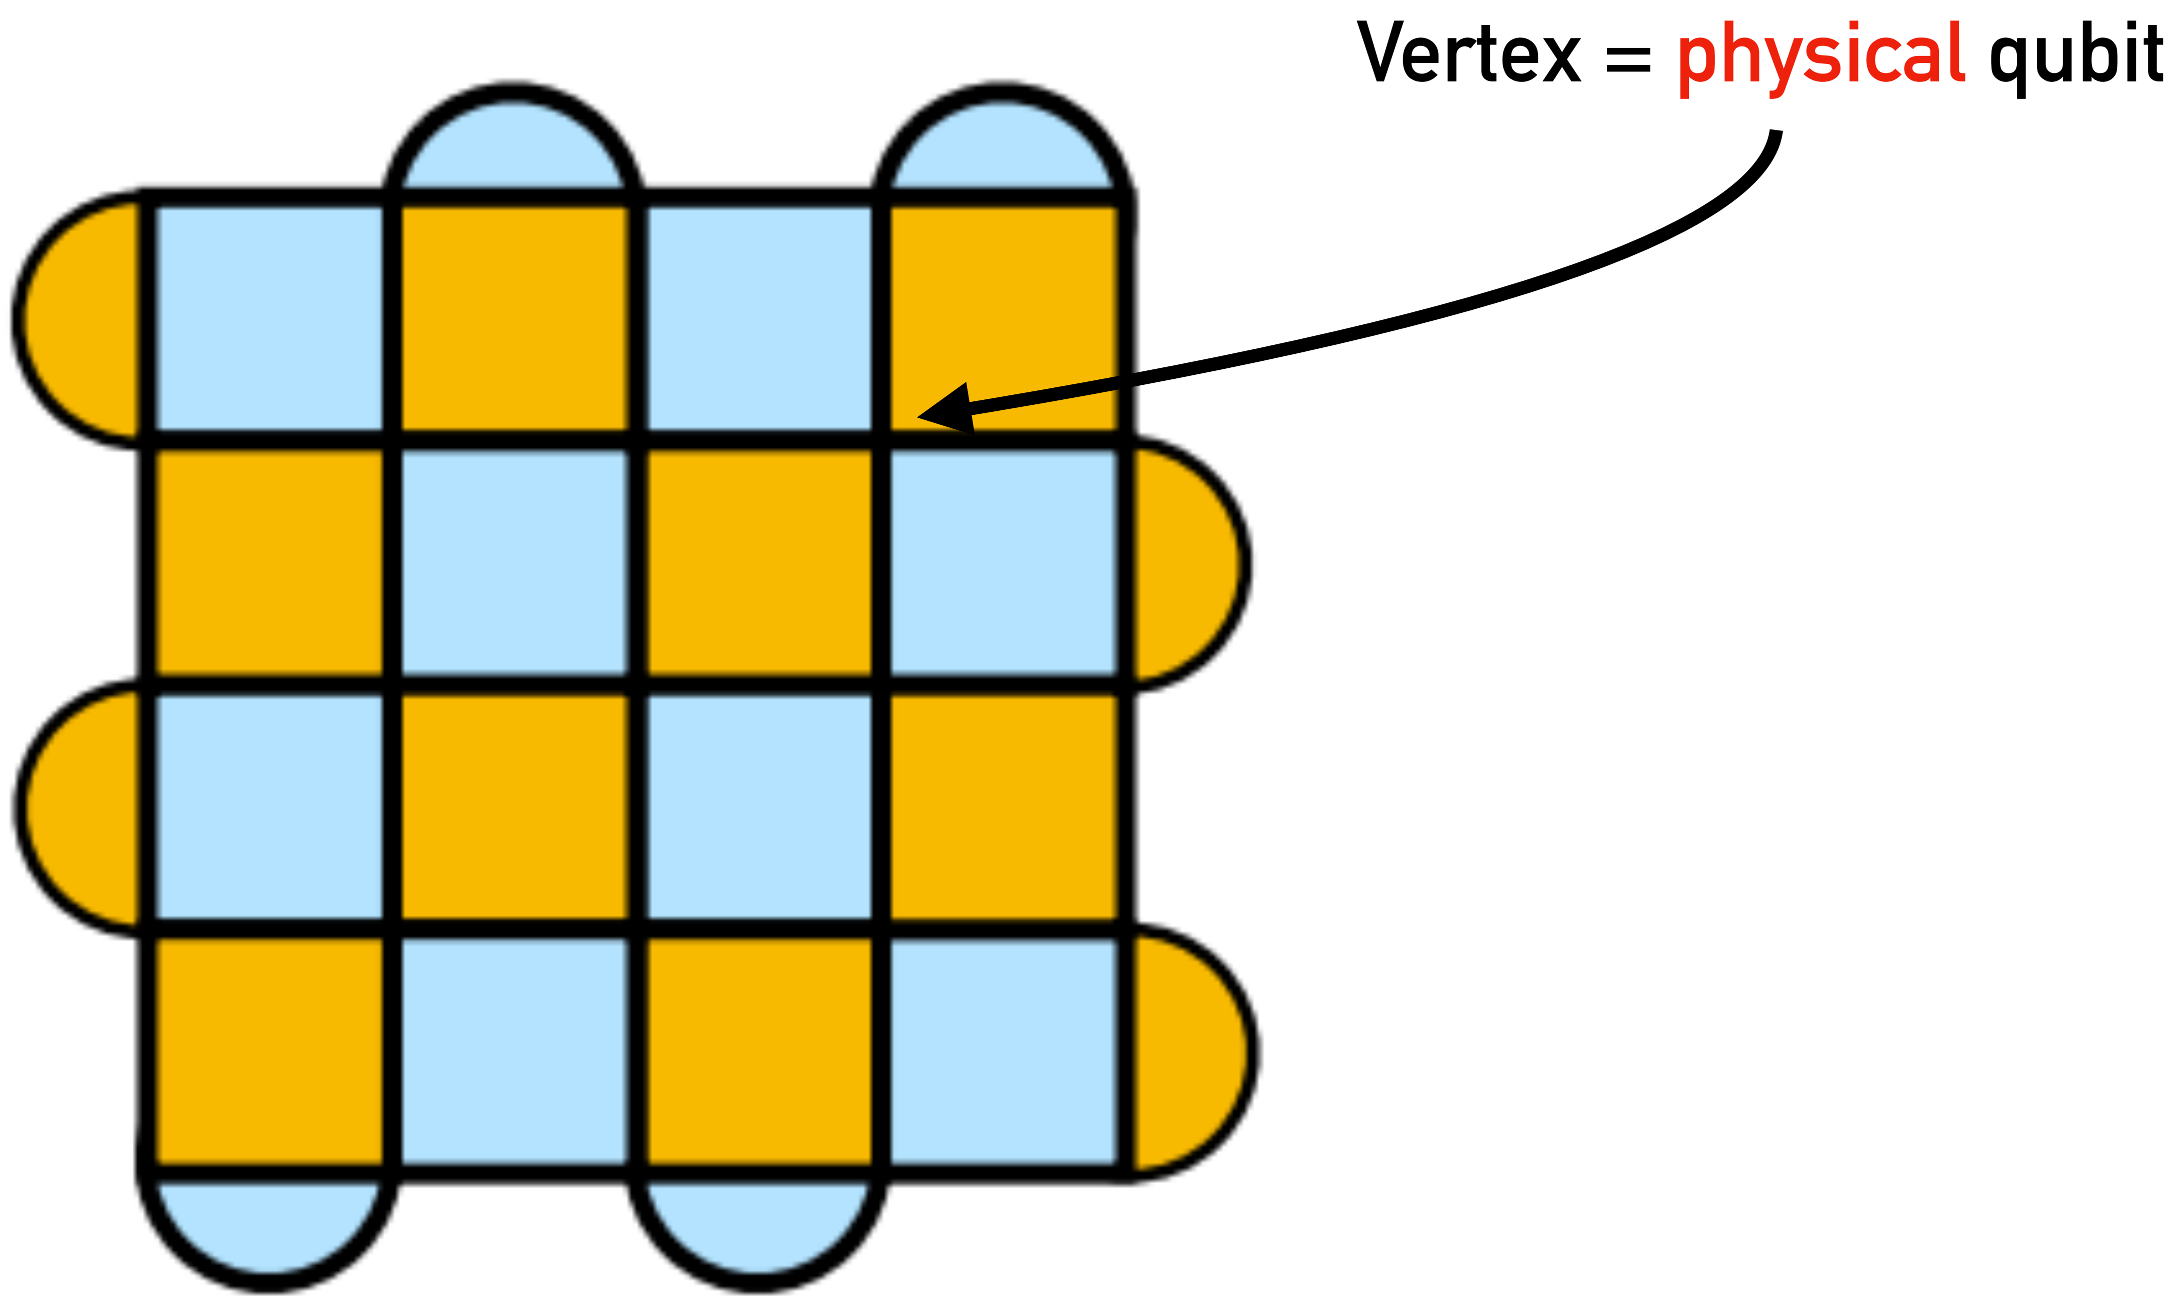

A simple quantum computer with one logical qubit

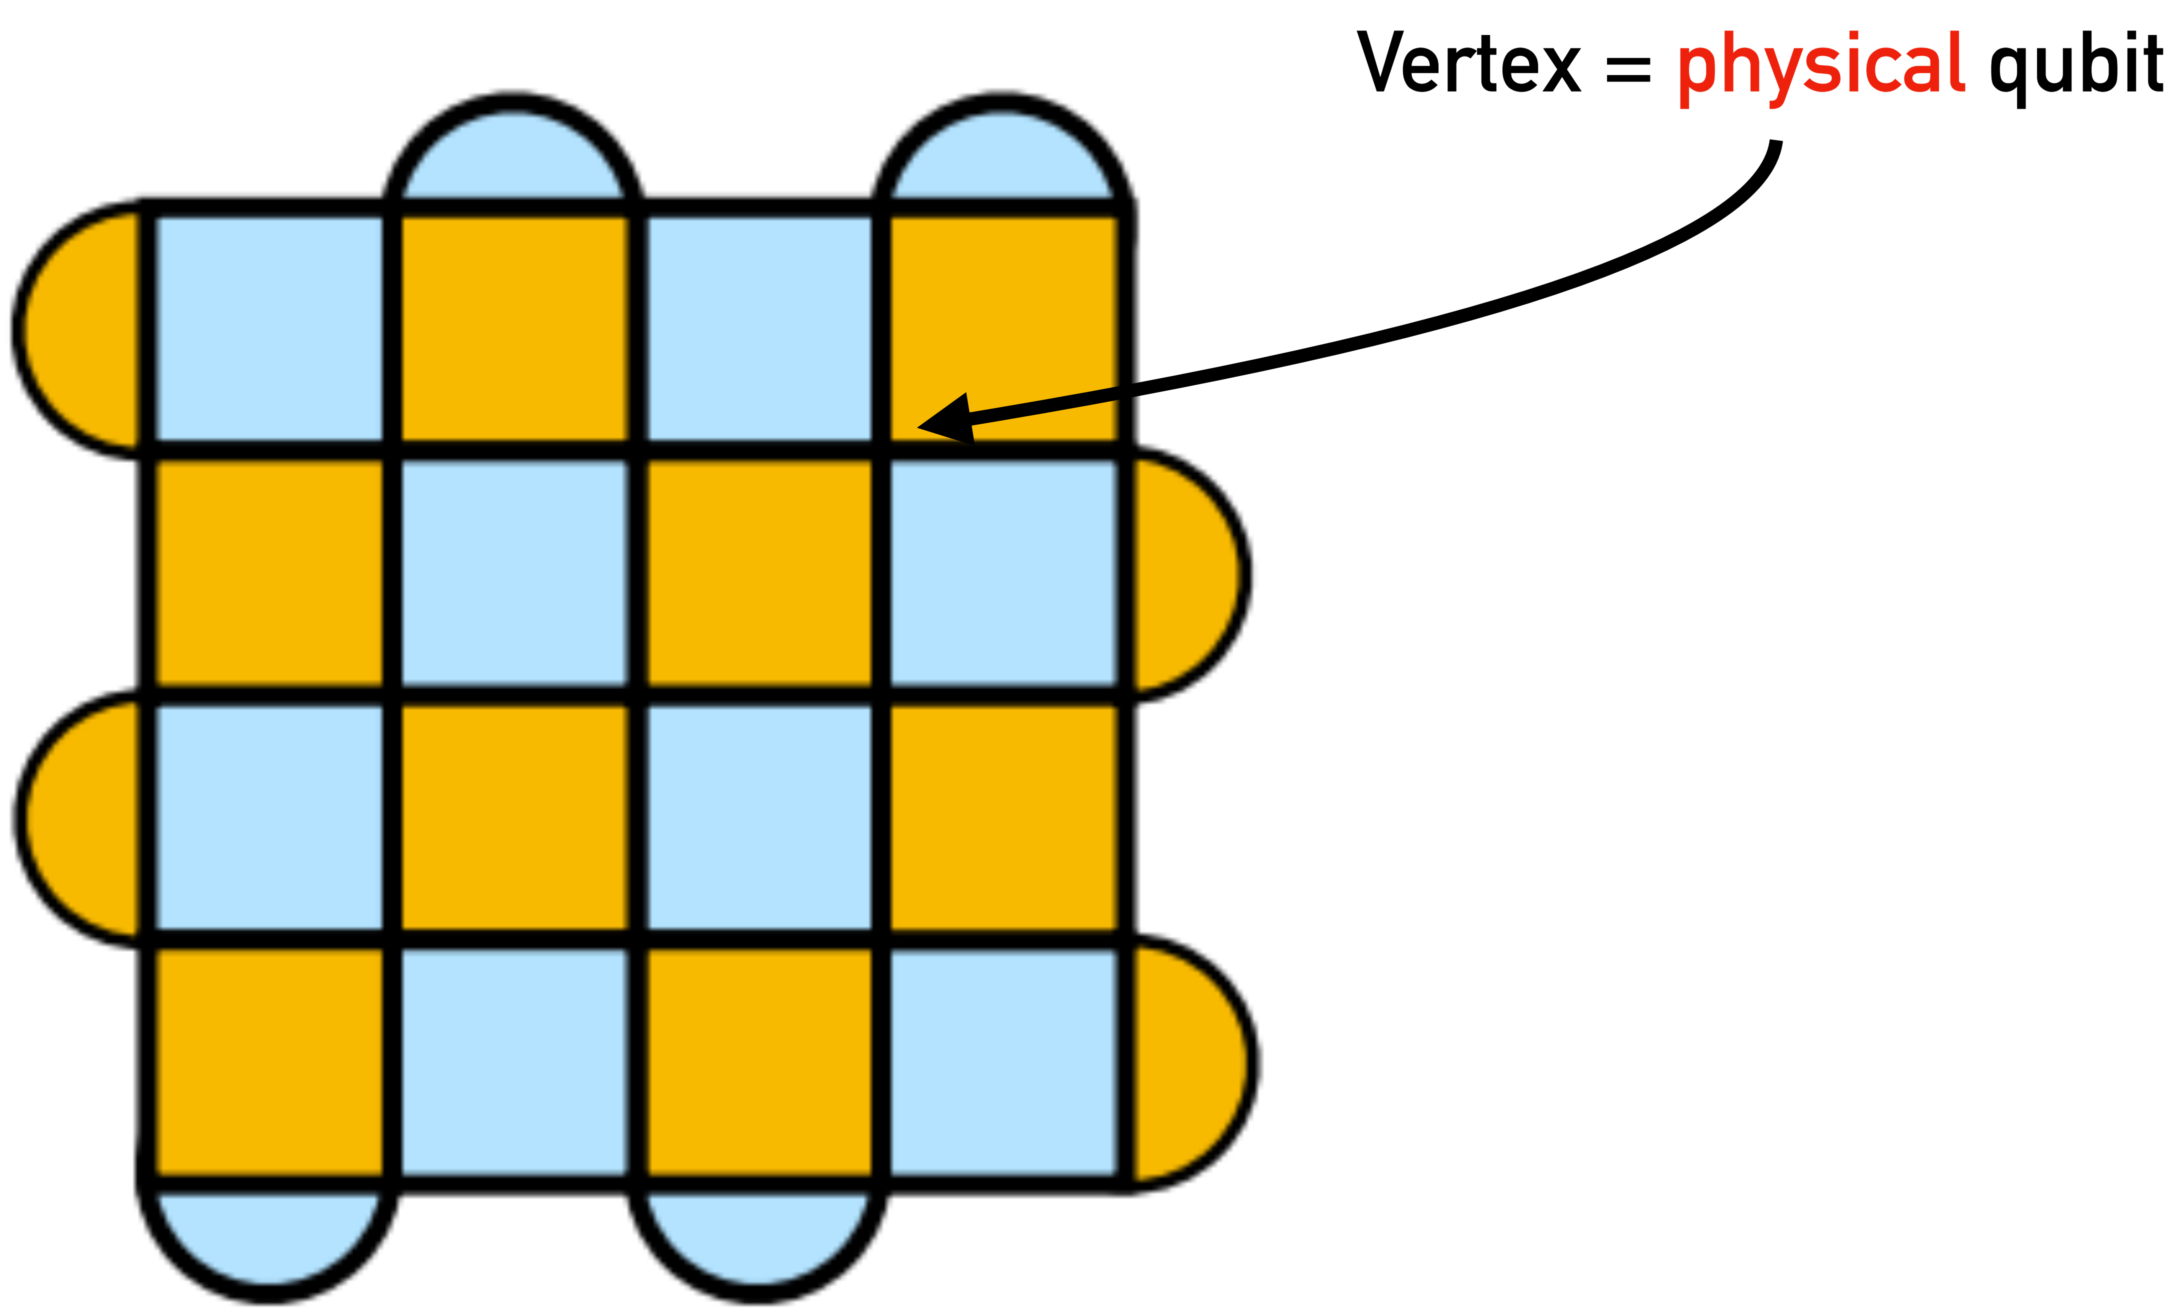

The full state of all these qubits itself represents a "logical" qubit

Simplification: all qubits 0 -> logical 0, all qubits 1 -> logical 1

Qubit errors show up as red dots on the orange squares

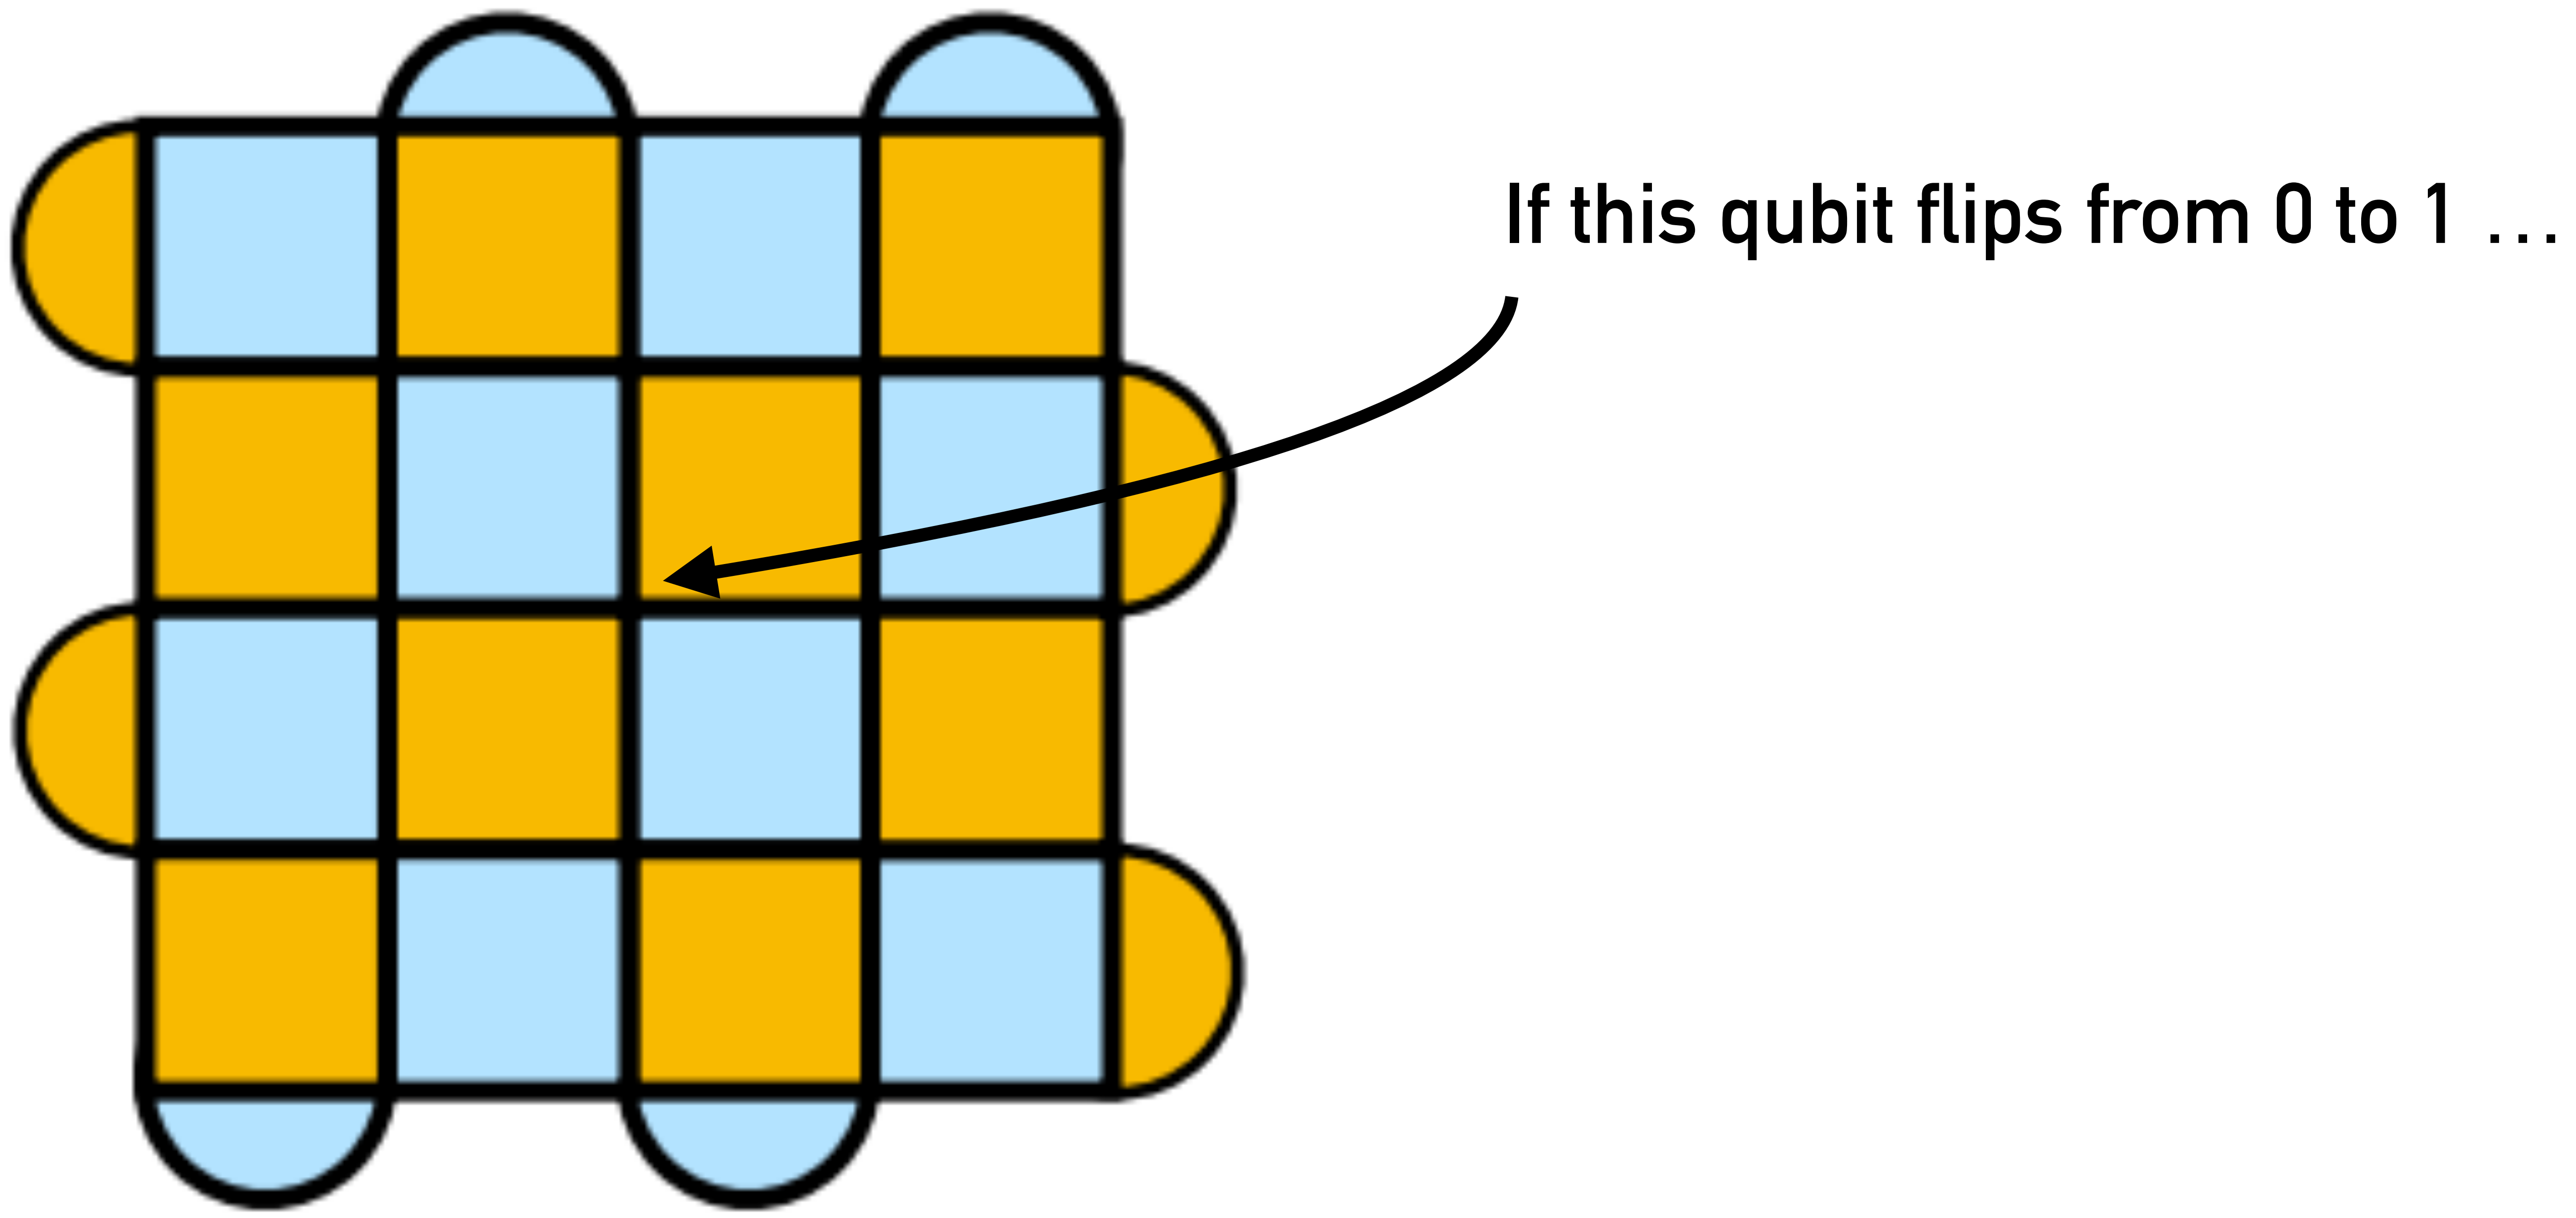

Qubit errors show up as red dots on the orange squares

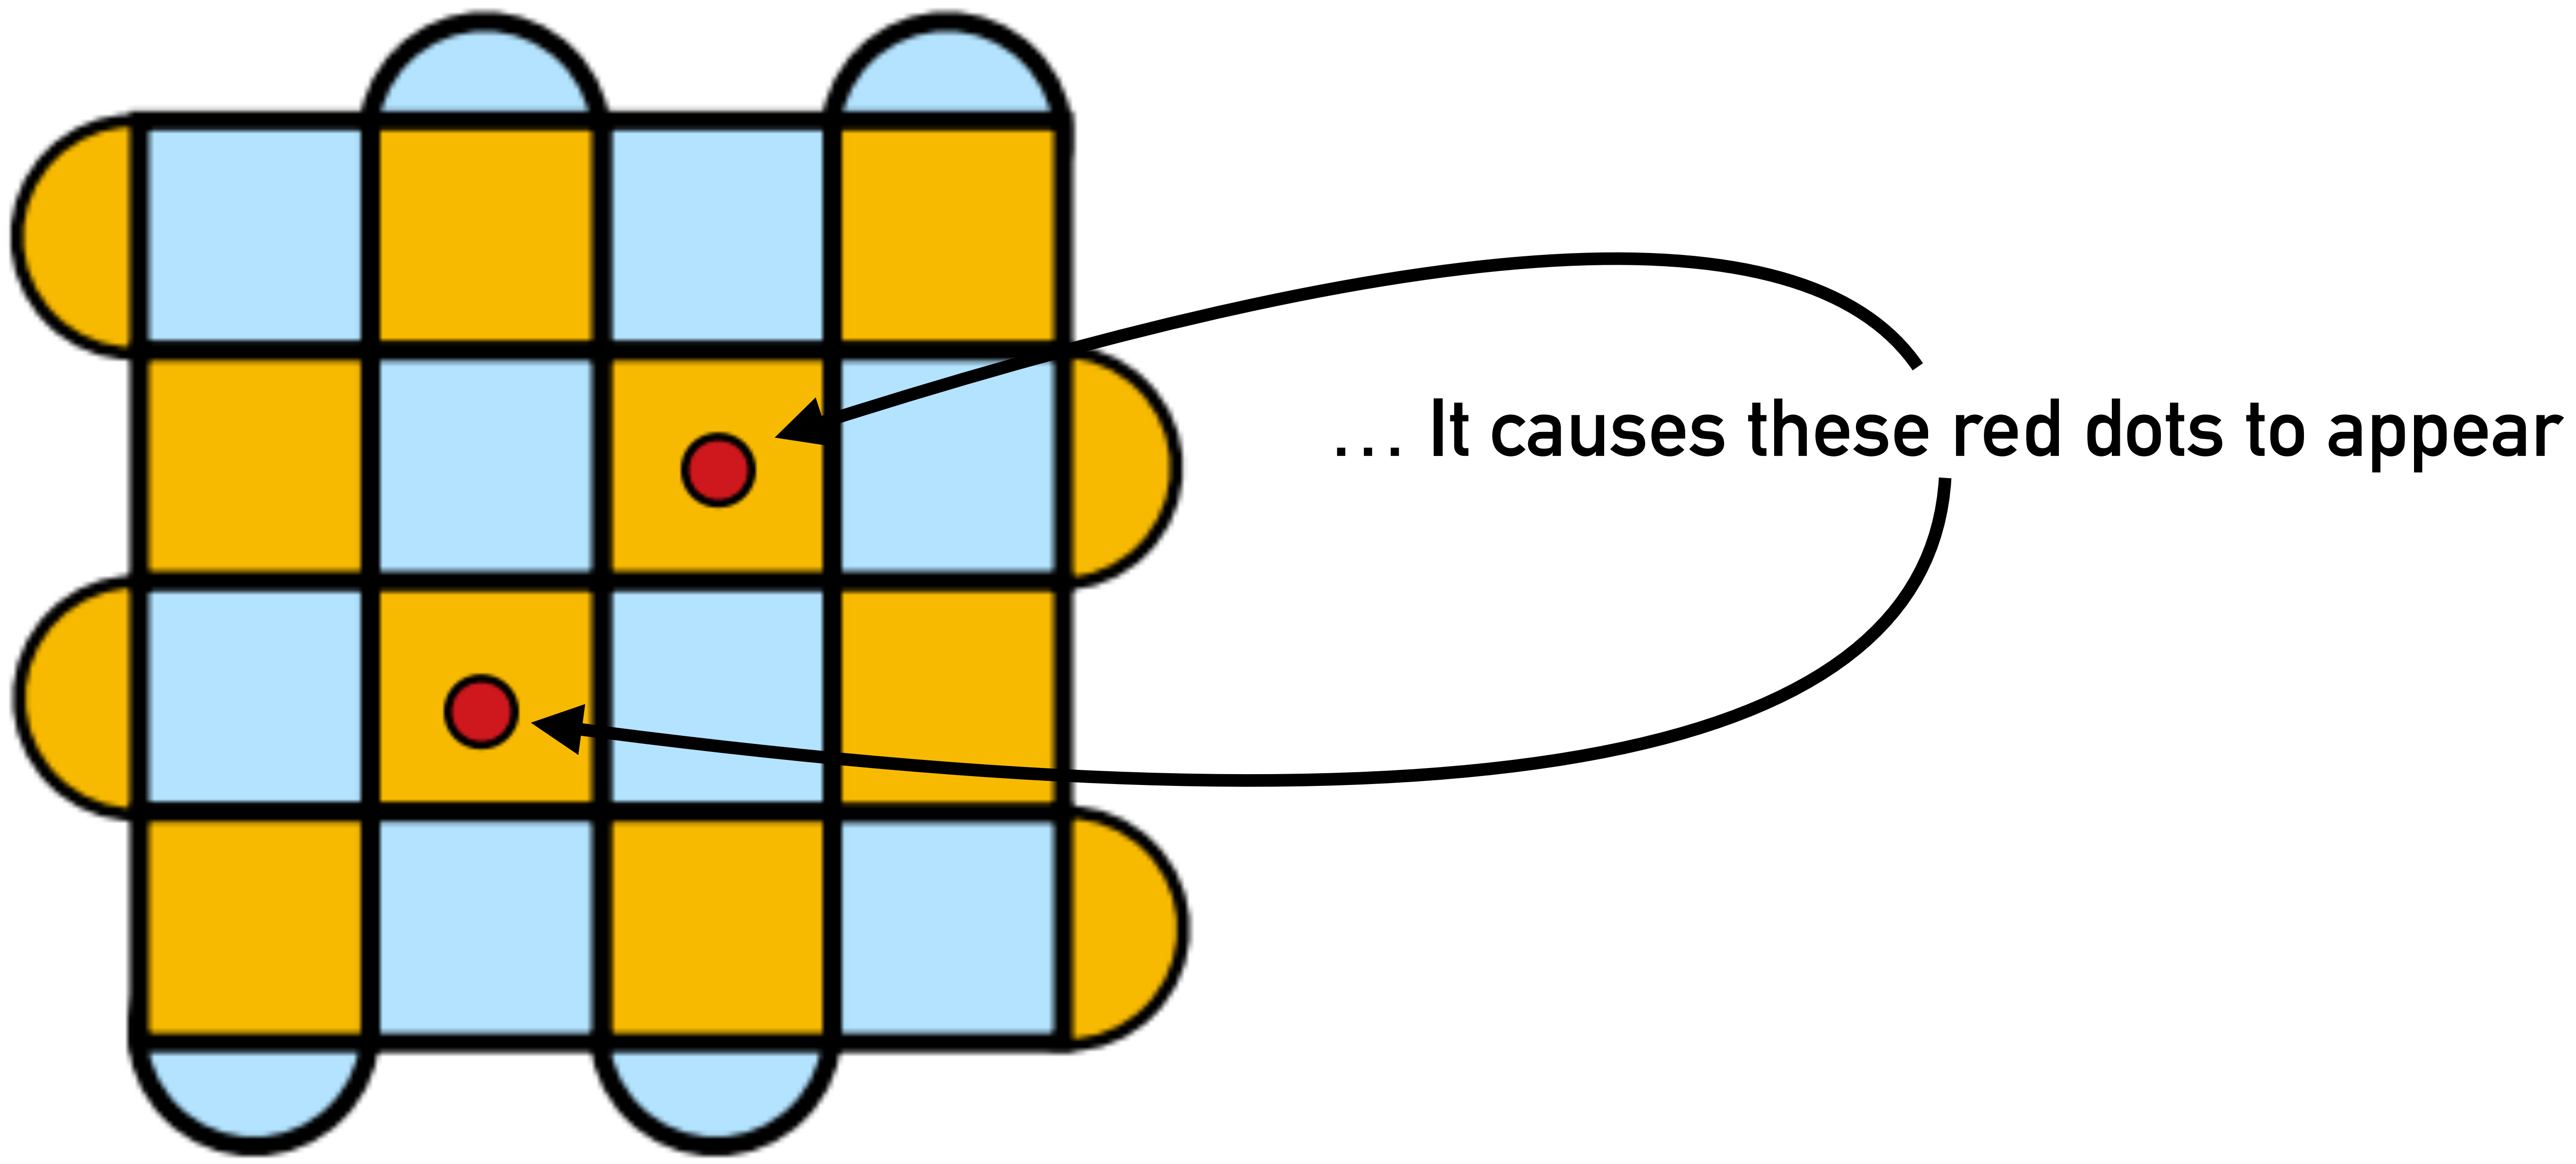

We can *not* look at the flipped qubits, only the red dots!

Multiple errors cause red dots to change position

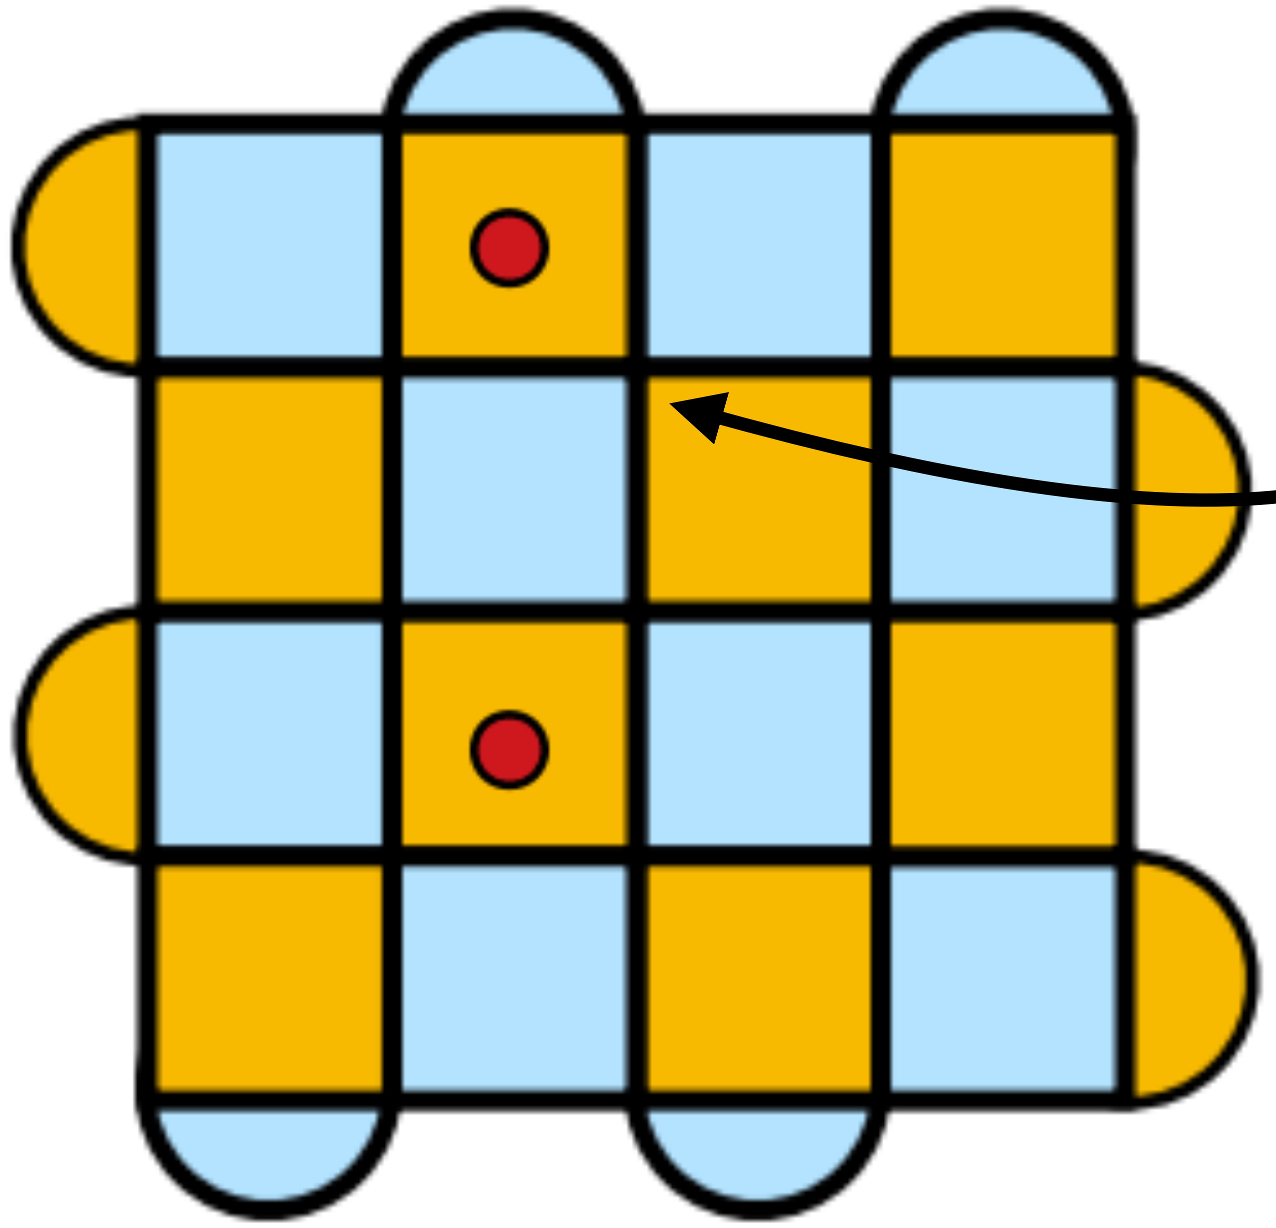

If this qubit flips from 0 to 1 too …

Only an odd number of red flags is visible

Red dots can (dis)appear at the edges

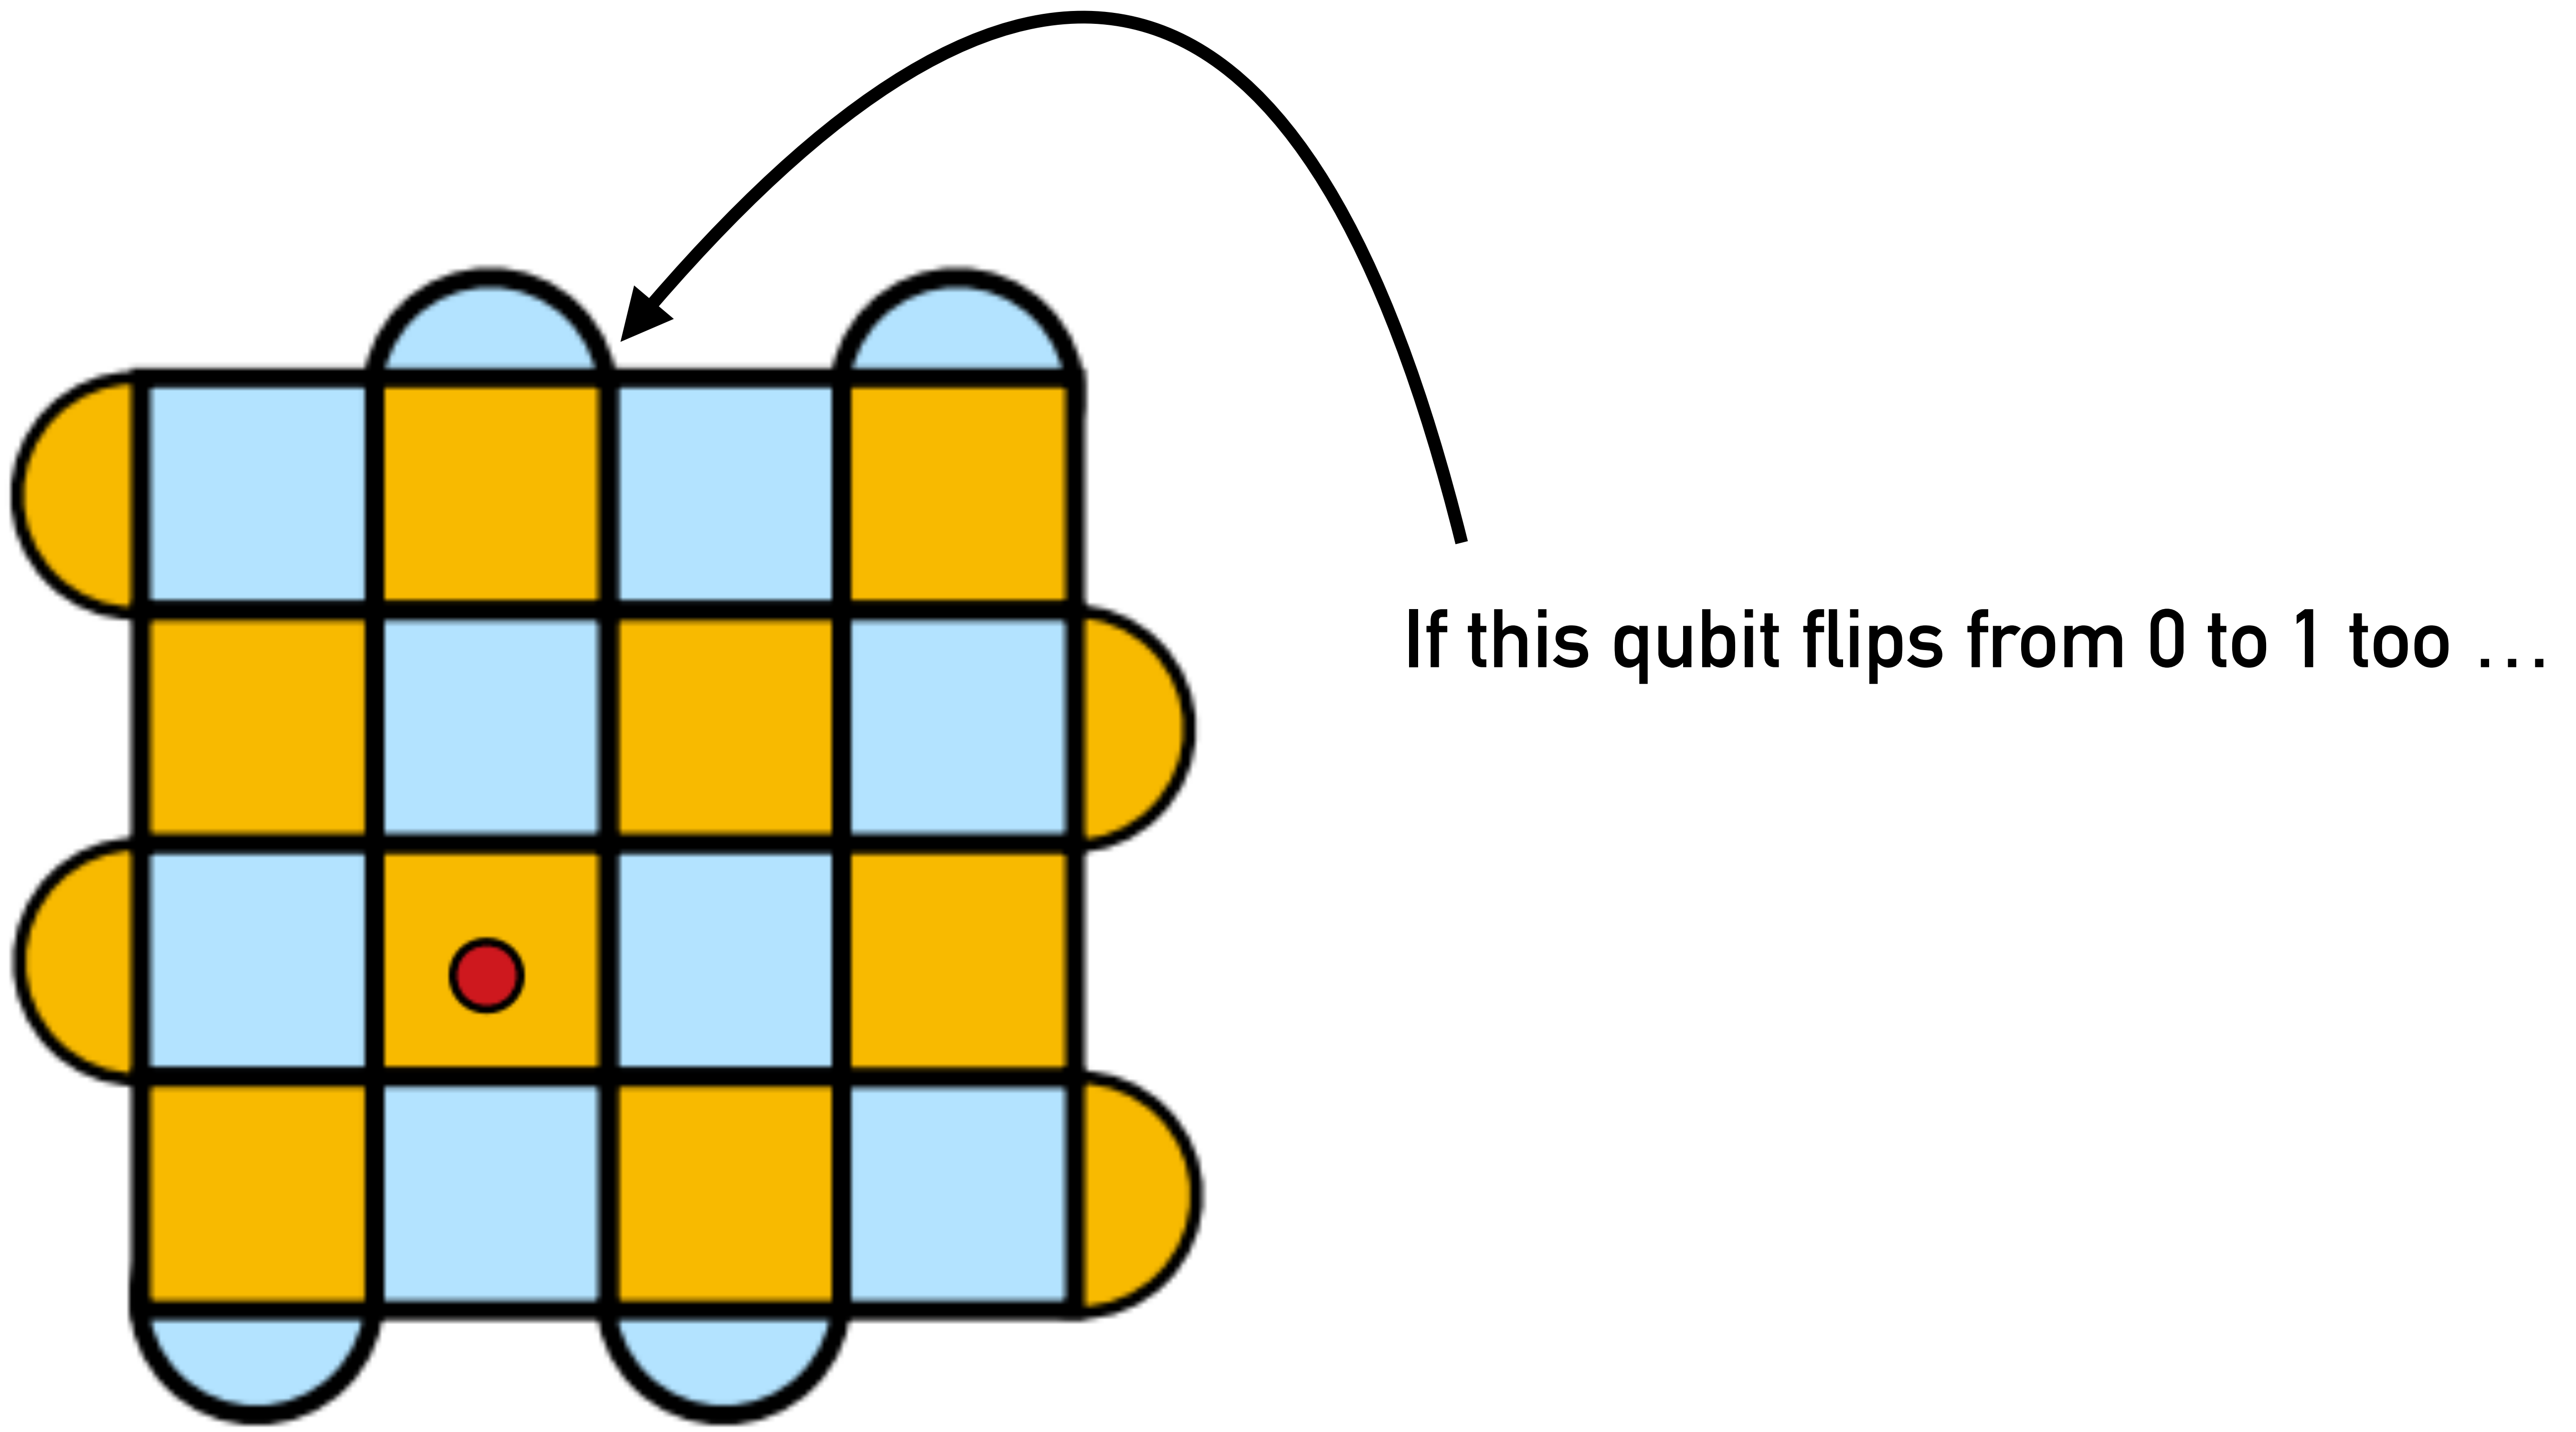

If a string of errors connects the edges, it is impossible to find out which qubits had errors

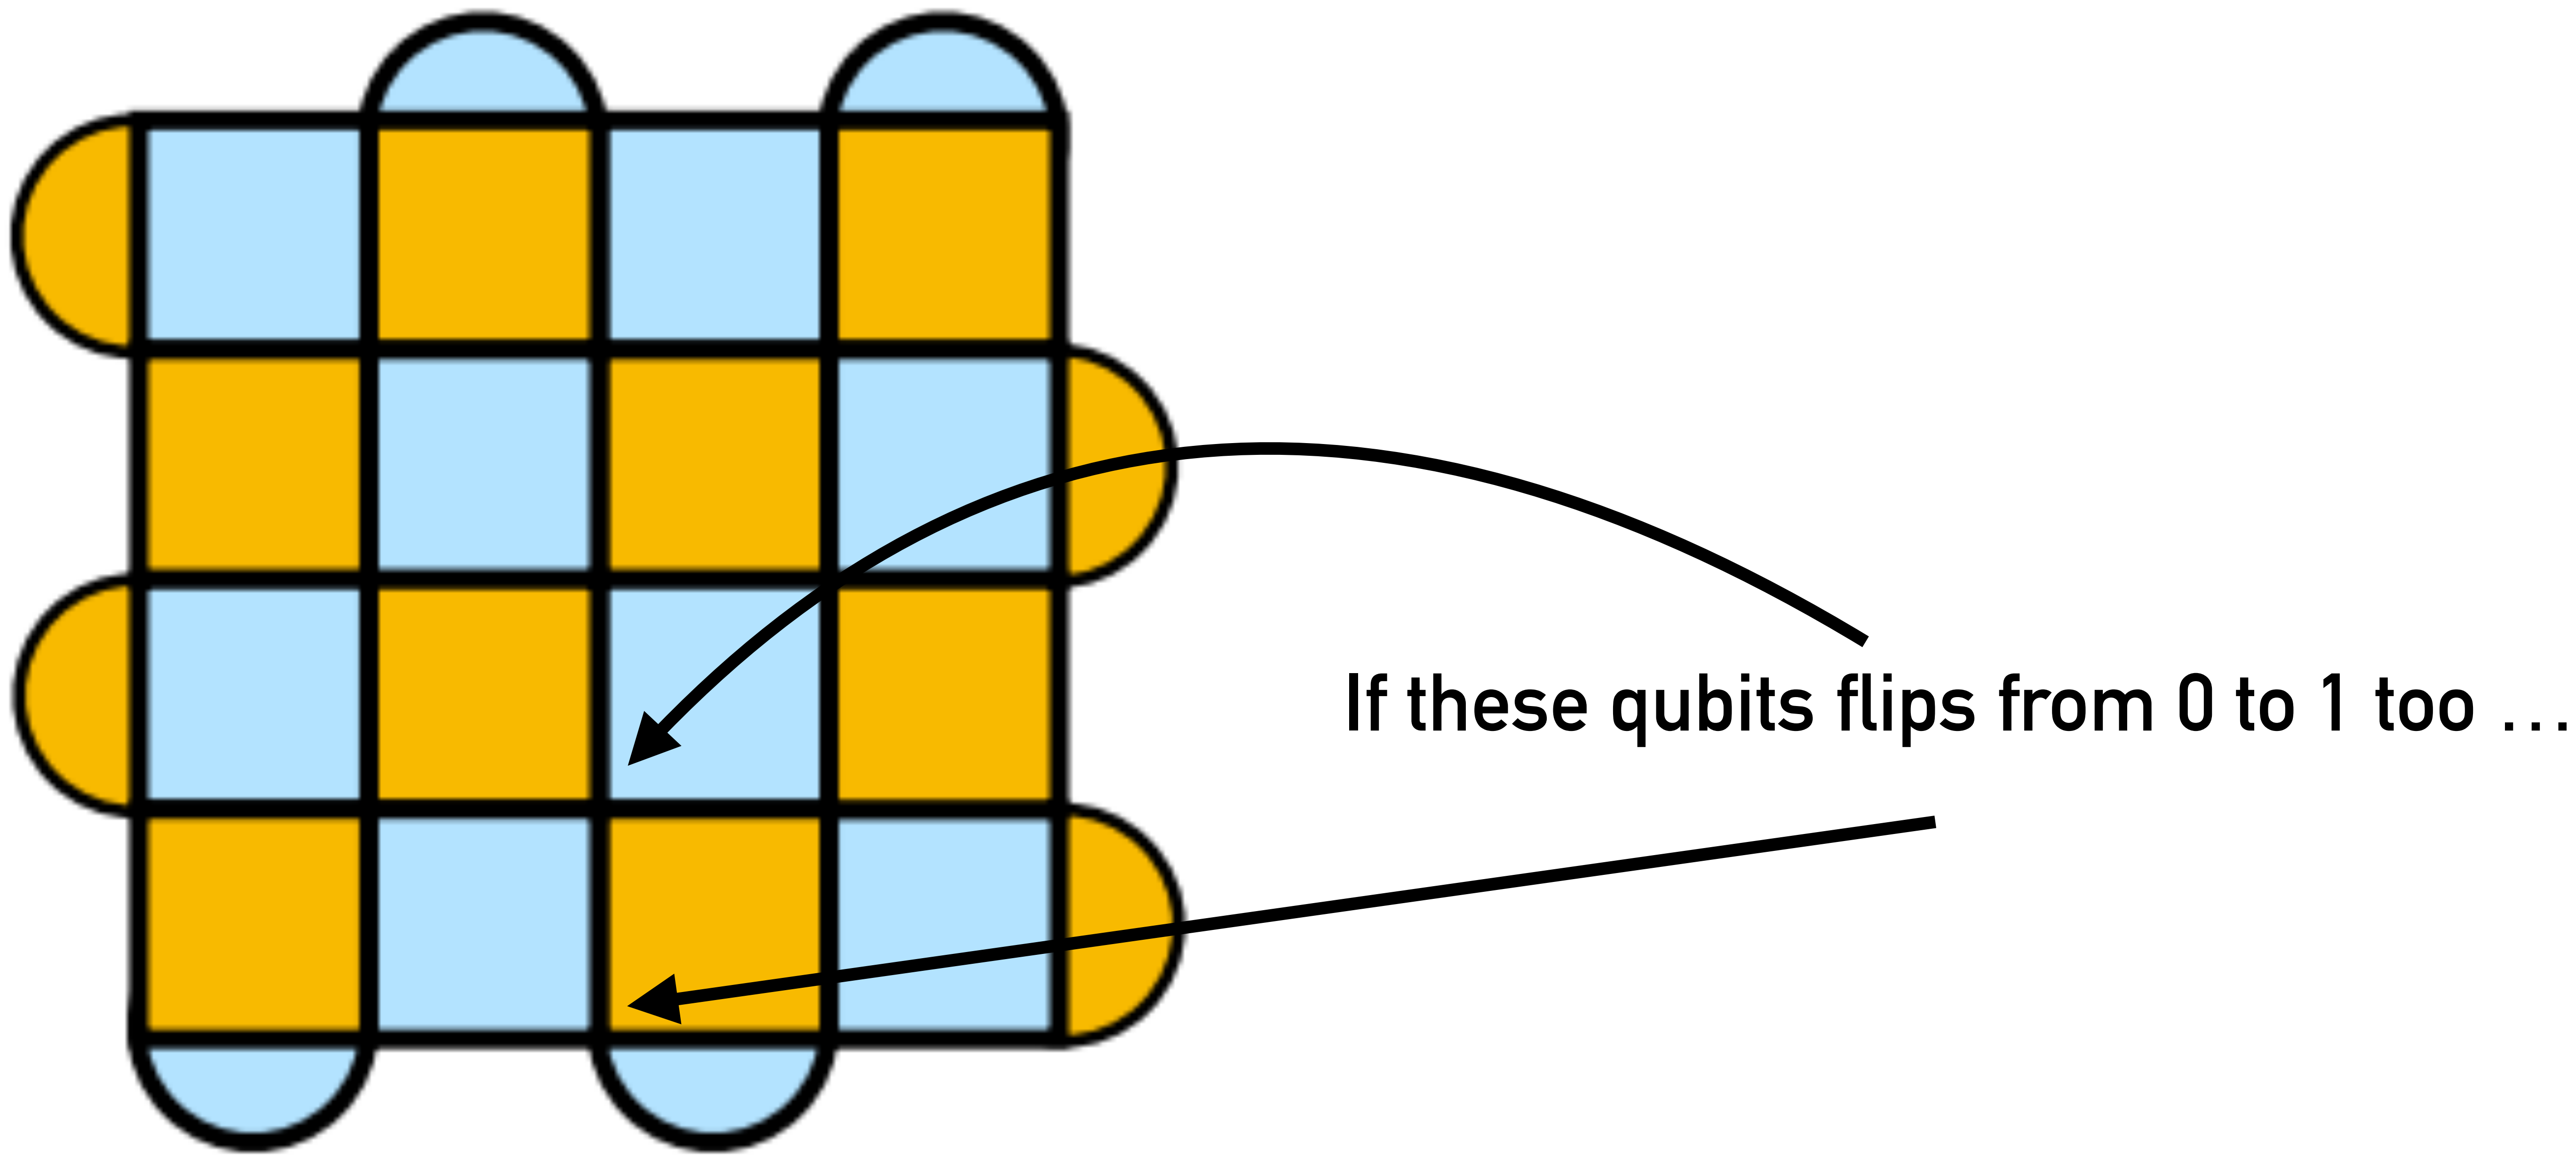

A whole column of qubits has errors, but we can't see it!

So quantum error correction is like a board game!

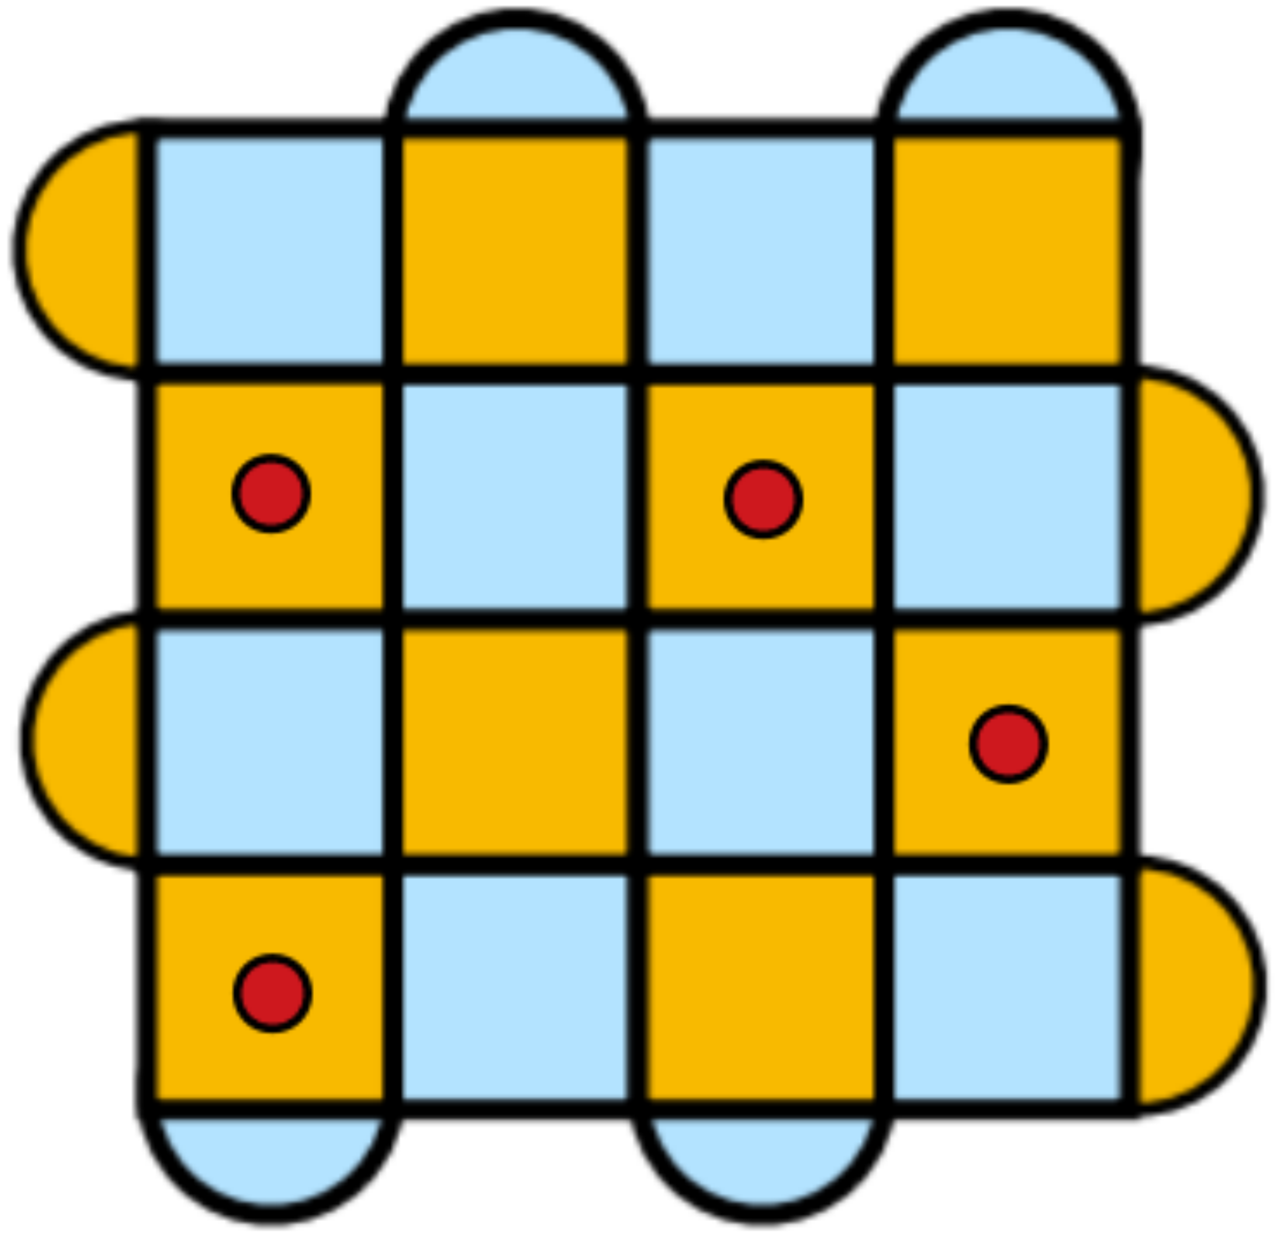

Given red dots, find out which qubits flipped (the errors)

Game over if a string of errors connects one edge to the other

Reinforcement learning can do quantum error correction

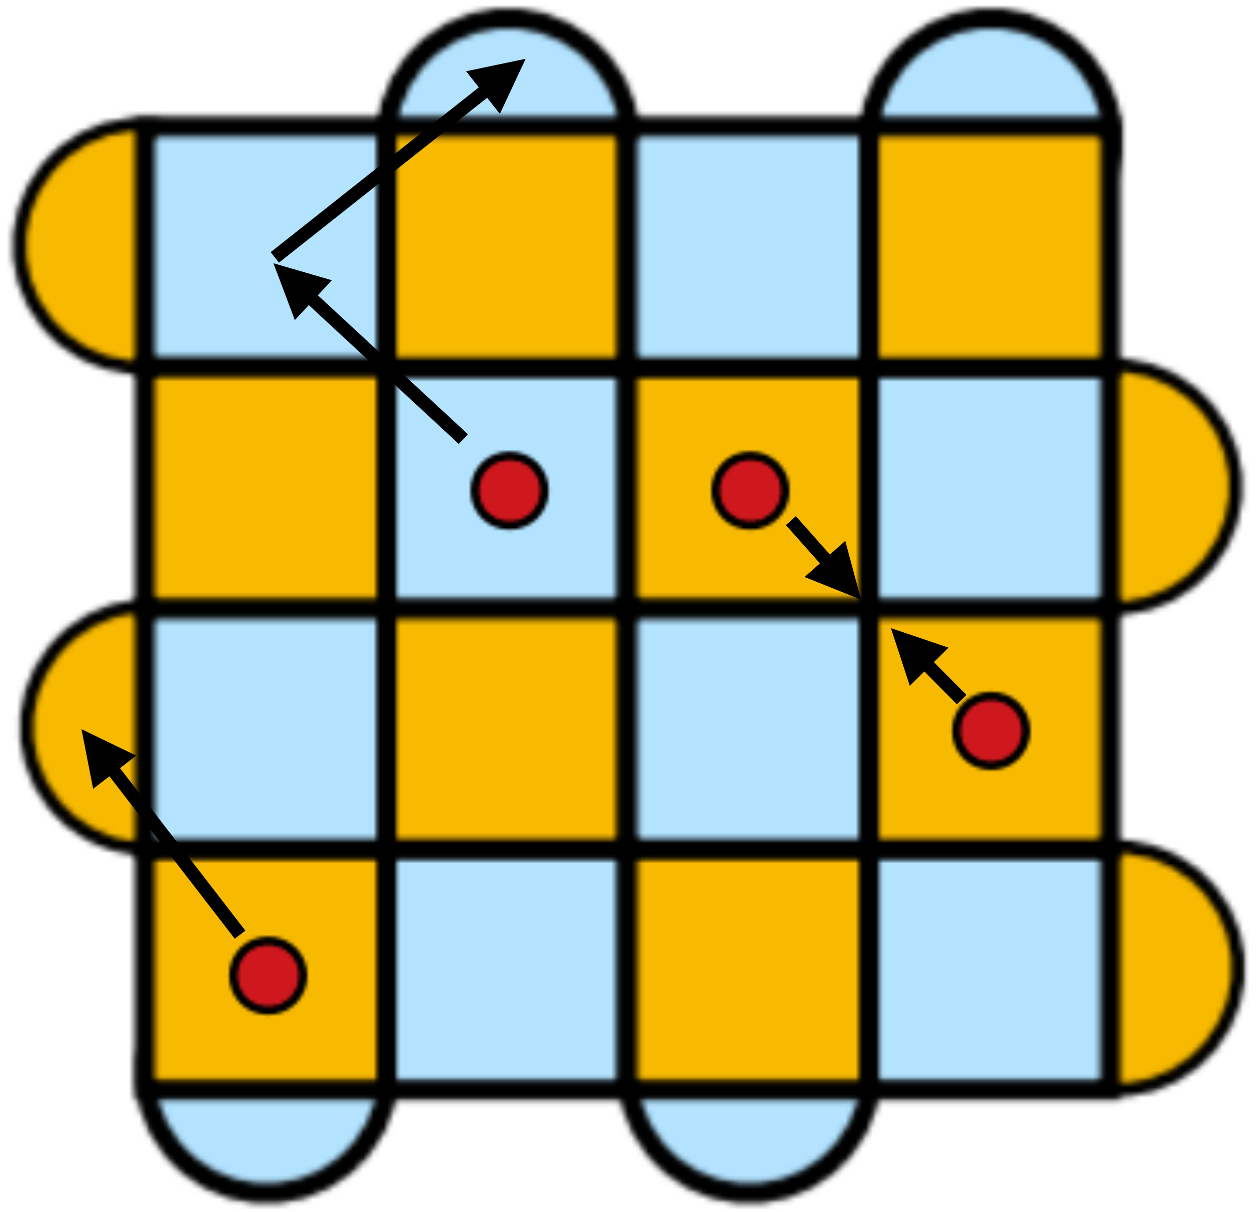

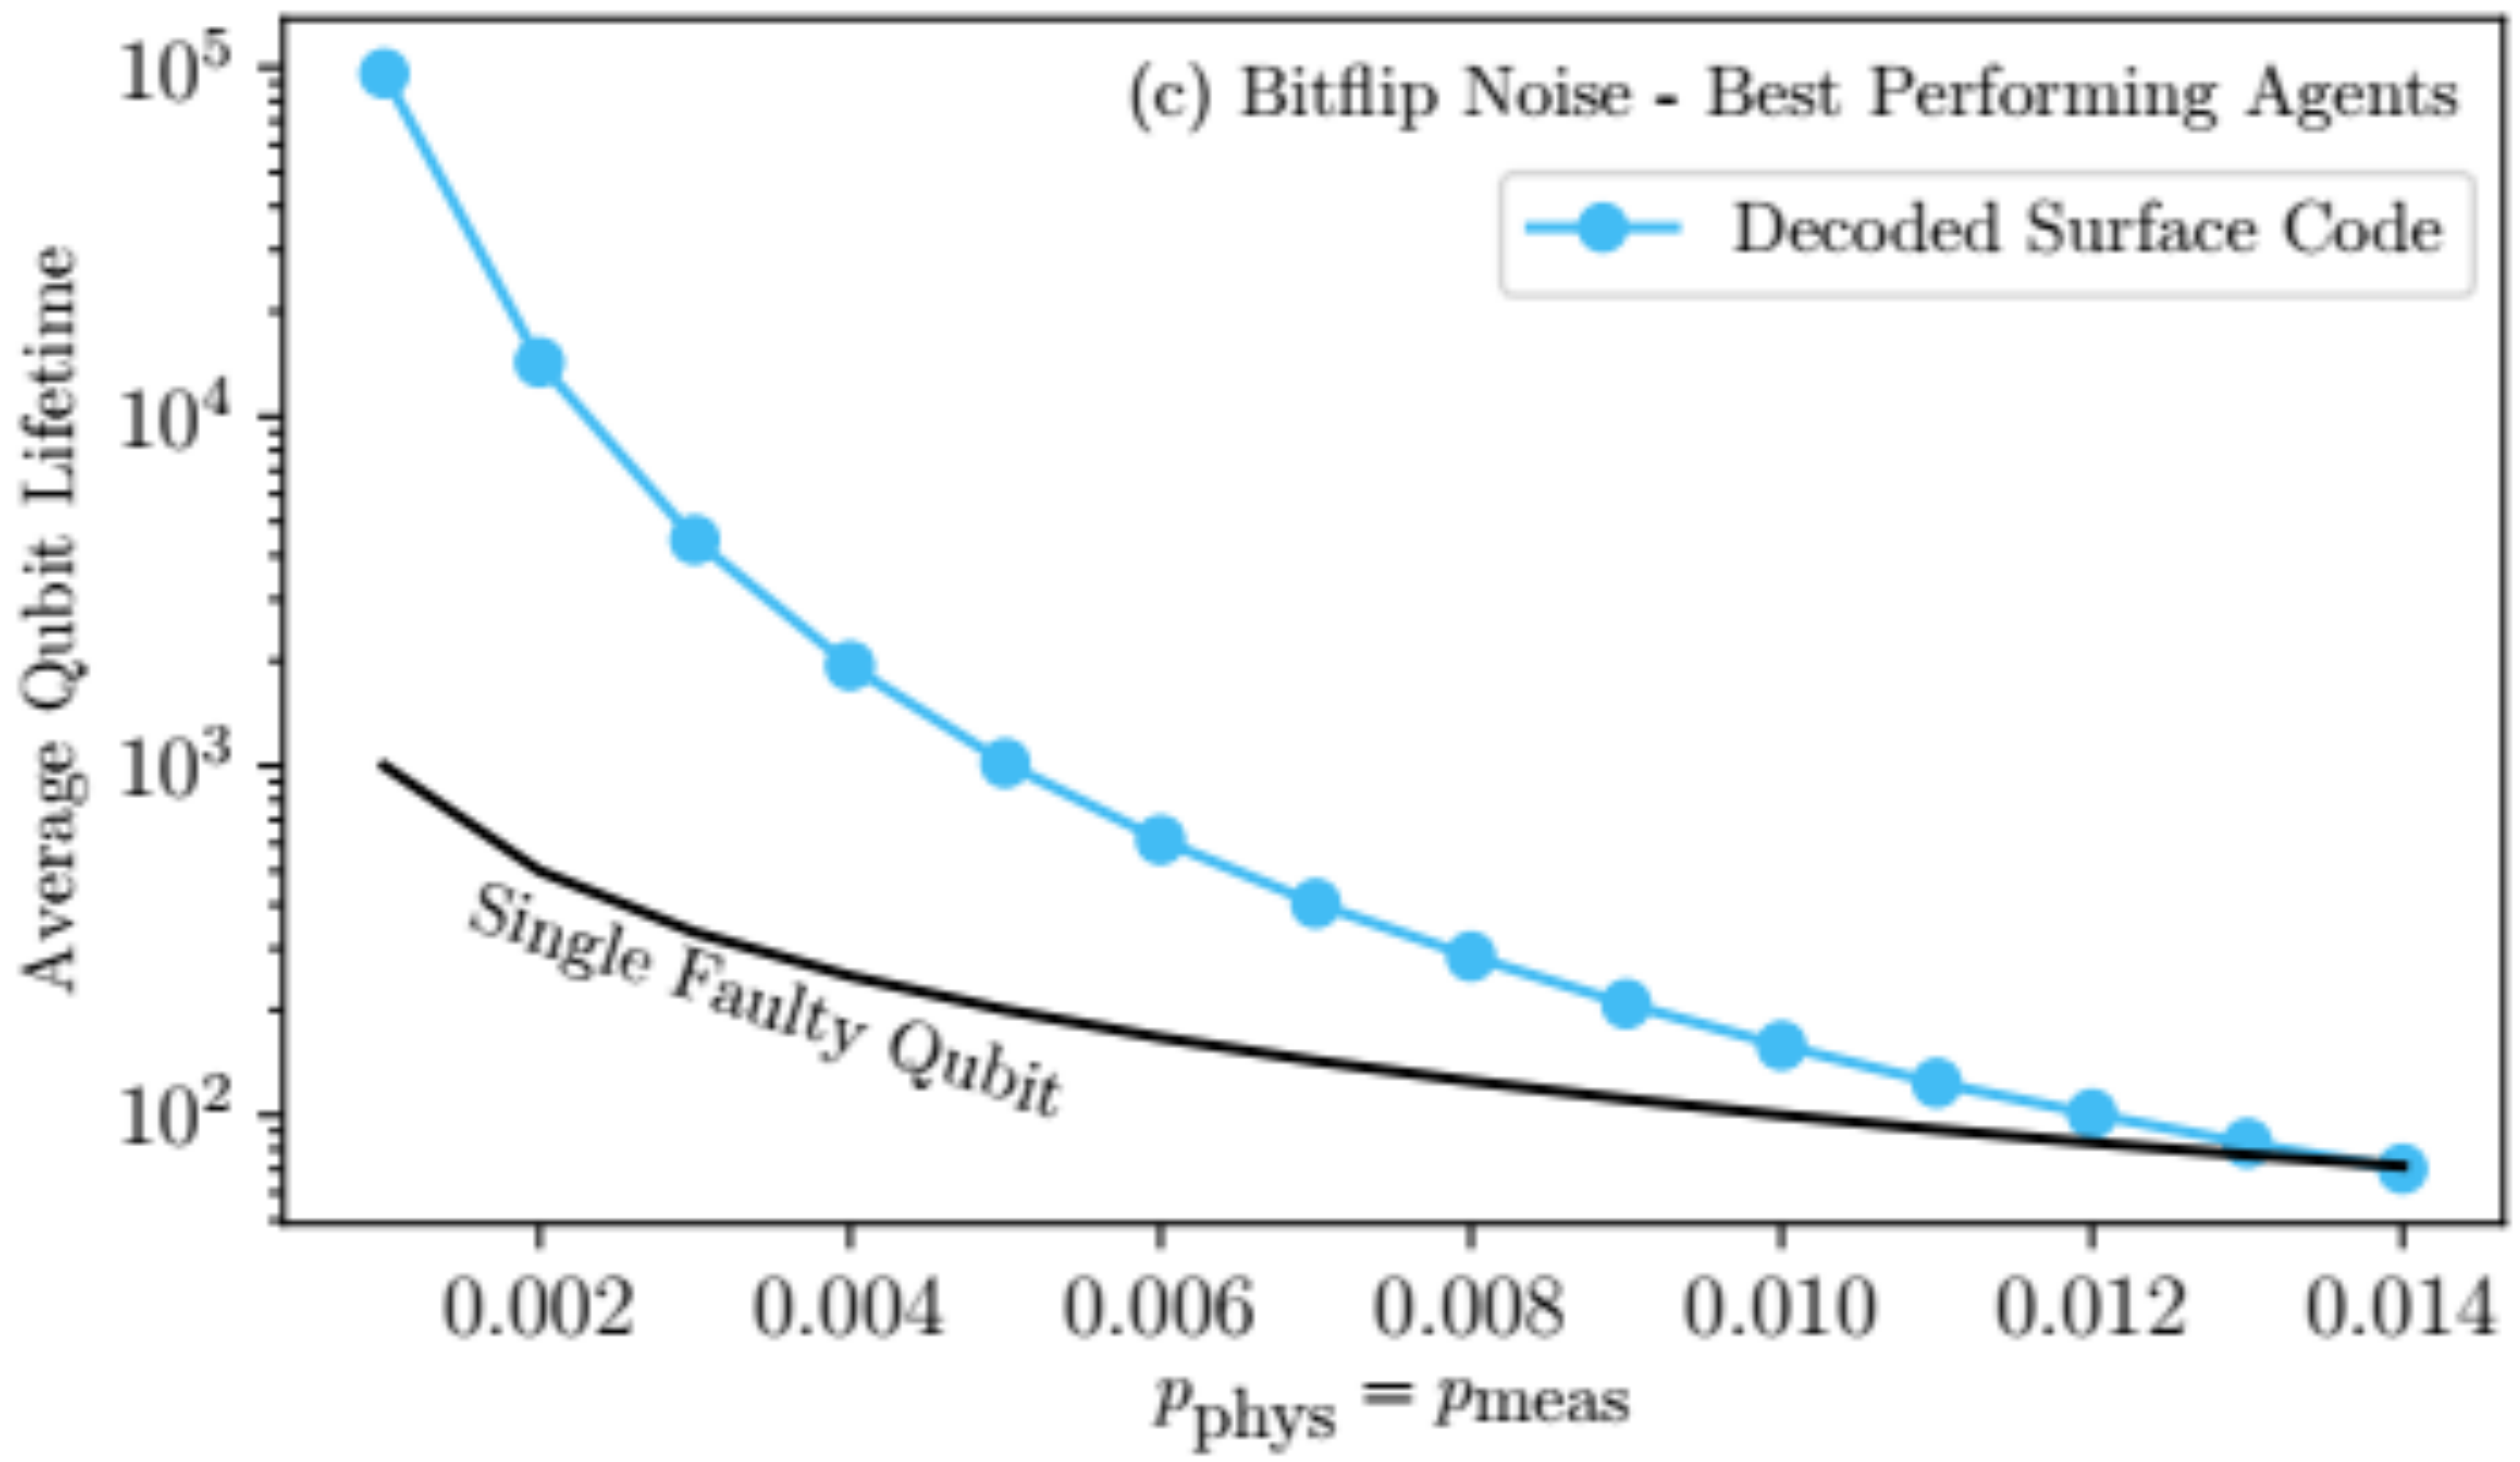

Use a neural network to determine which qubits flipped, given the current red dots

These methods are not mutually exclusive

A glimpse of the future?

Control experiments

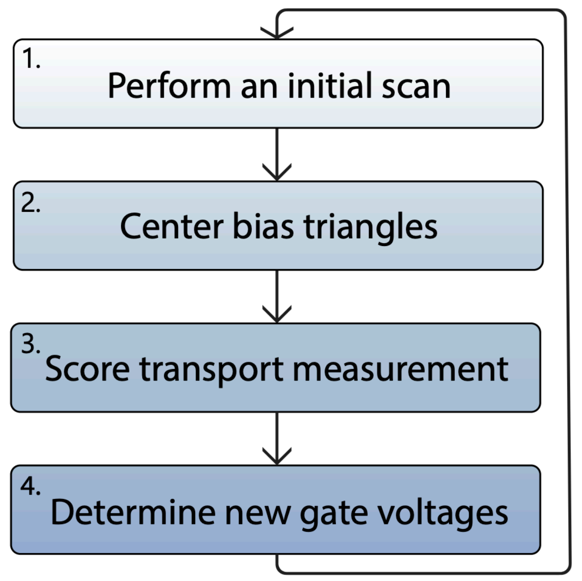

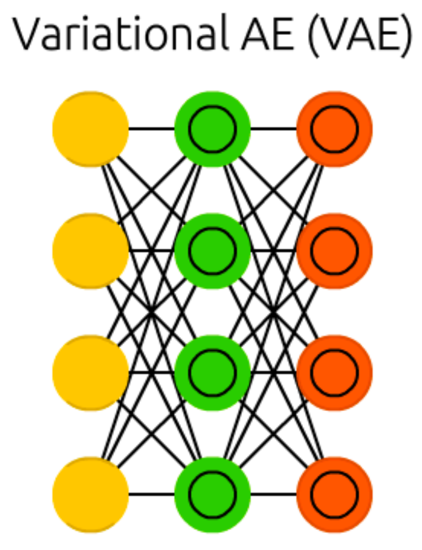

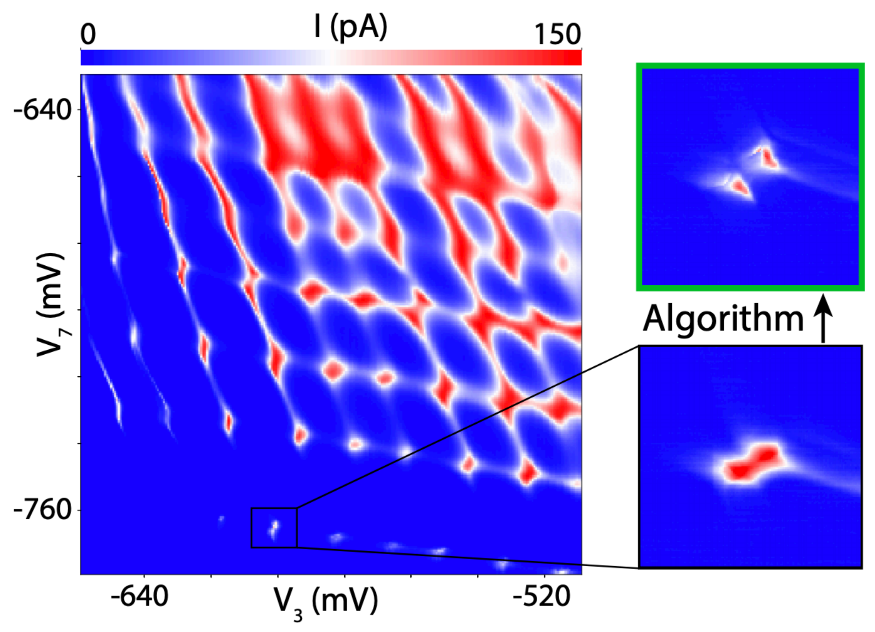

Natalia Ares' Lab, arXiv:2001.04409

Design experiments

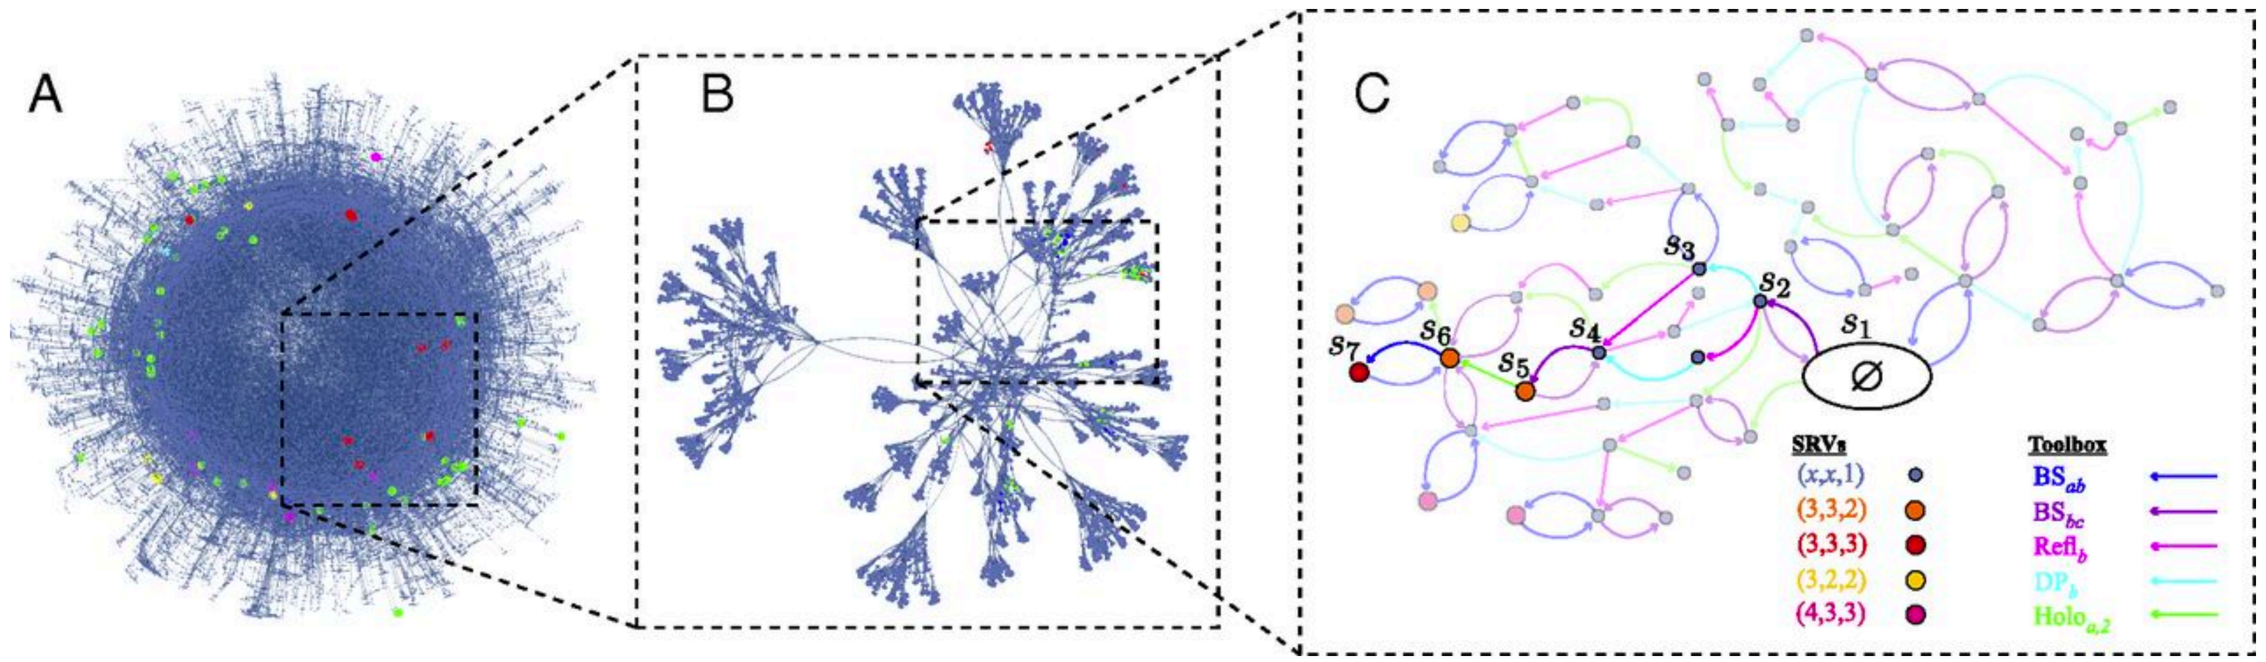

Melnikov et al, PNAS February 6, 2018 115 (6) 1221-1226

When does an ML approach make sense in physics?

Accuracy

Only ever as accurate as the data (unless ML can also request new data?)

Speed
Experimental/Computational Cost

Predict ground state properties vs Monte Carlo, or learn a quantum state

quantummanybody.webflow.io

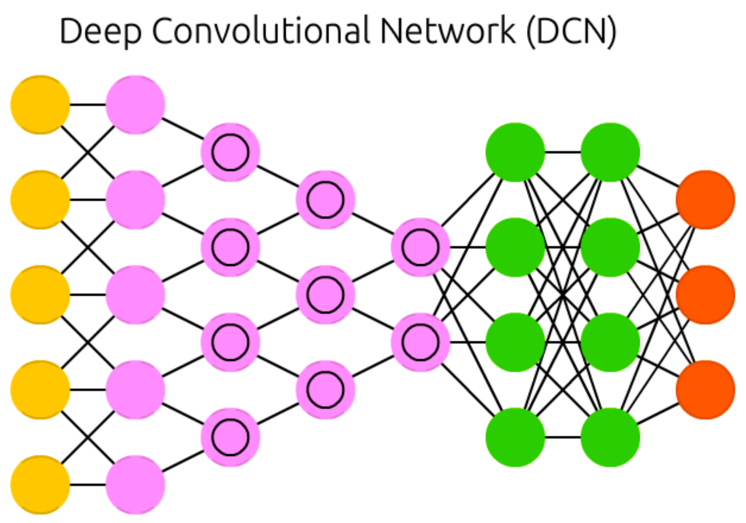

Adaptiveness

If qubit number 37 happens to be worse than the others, ML will learn that

Quantum Games

https://arxiv.org/abs/1801.00862

Quantum Computing in the NISQ era and beyond

John Preskill

6.11 Quantum games

Advances in classical computing launched a new world of digital games, touching the lives of millions and generating billions in revenue. Could quantum computers do the same?

Physicists often say that the quantum world is counter-intuitive because it is so foreign to ordinary experience. That's true now, but might it be different in the future? Perhaps kids who grow up playing quantum games will acquire a visceral understanding of quantum phenomena that our generation lacks. Furthermore, quantum games could open a niche for quantum machine learning methods, which might seize the opportunity to improve game play in situations where quantum entanglement has an essential role.

Quantum Games

Quantum TiqTaqToe

www.quantumtictactoe.com

https://quantumfrontiers.com/2019/07/15/tiqtaqtoe/

www.quantumchess.net

Quantum Games

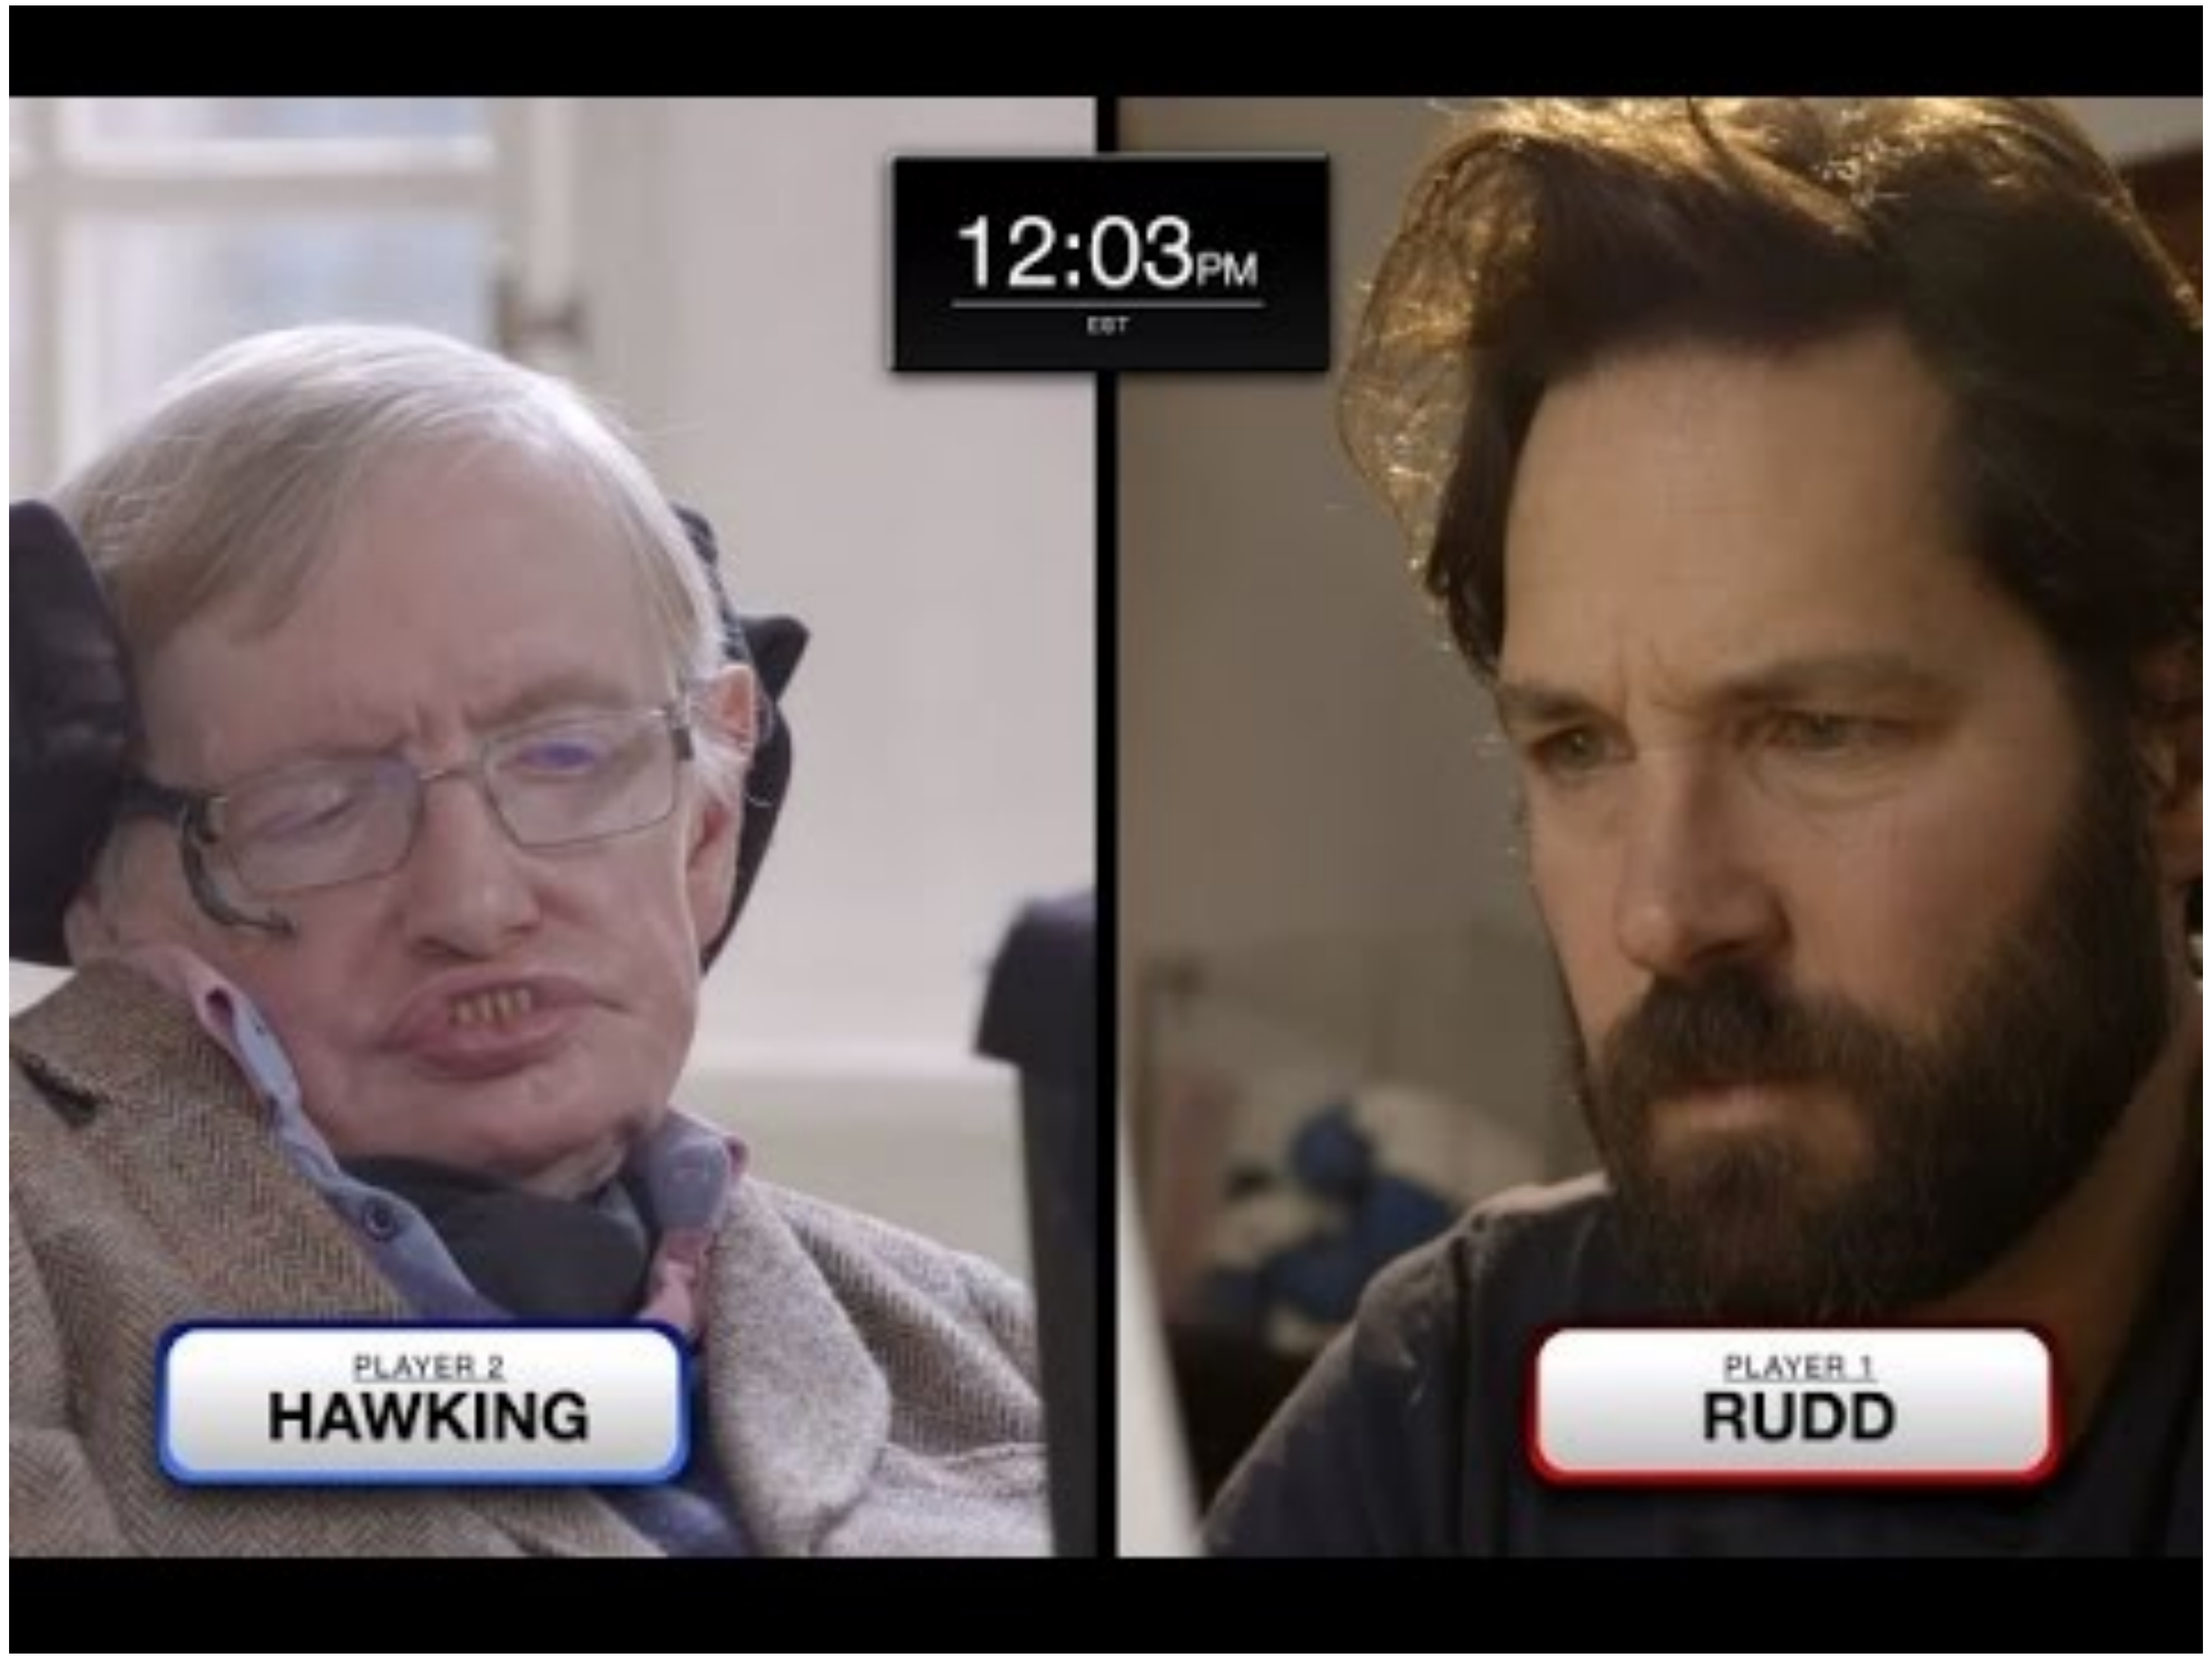

"Anyone can Quantum"

Thank you for participating!