Classical Machine Learning for Quantum Technologies

Florian Marquardt

Max Planck Institute for the Science of Light & Friedrich-Alexander-Universität, Erlangen/Germany

Introduction: Neural Networks
Overview: Machine Learning for Quantum Technologies
Deep Reinforcement Learning

Two examples:
Quantum Error Correction
Quantum Circuit Optimization

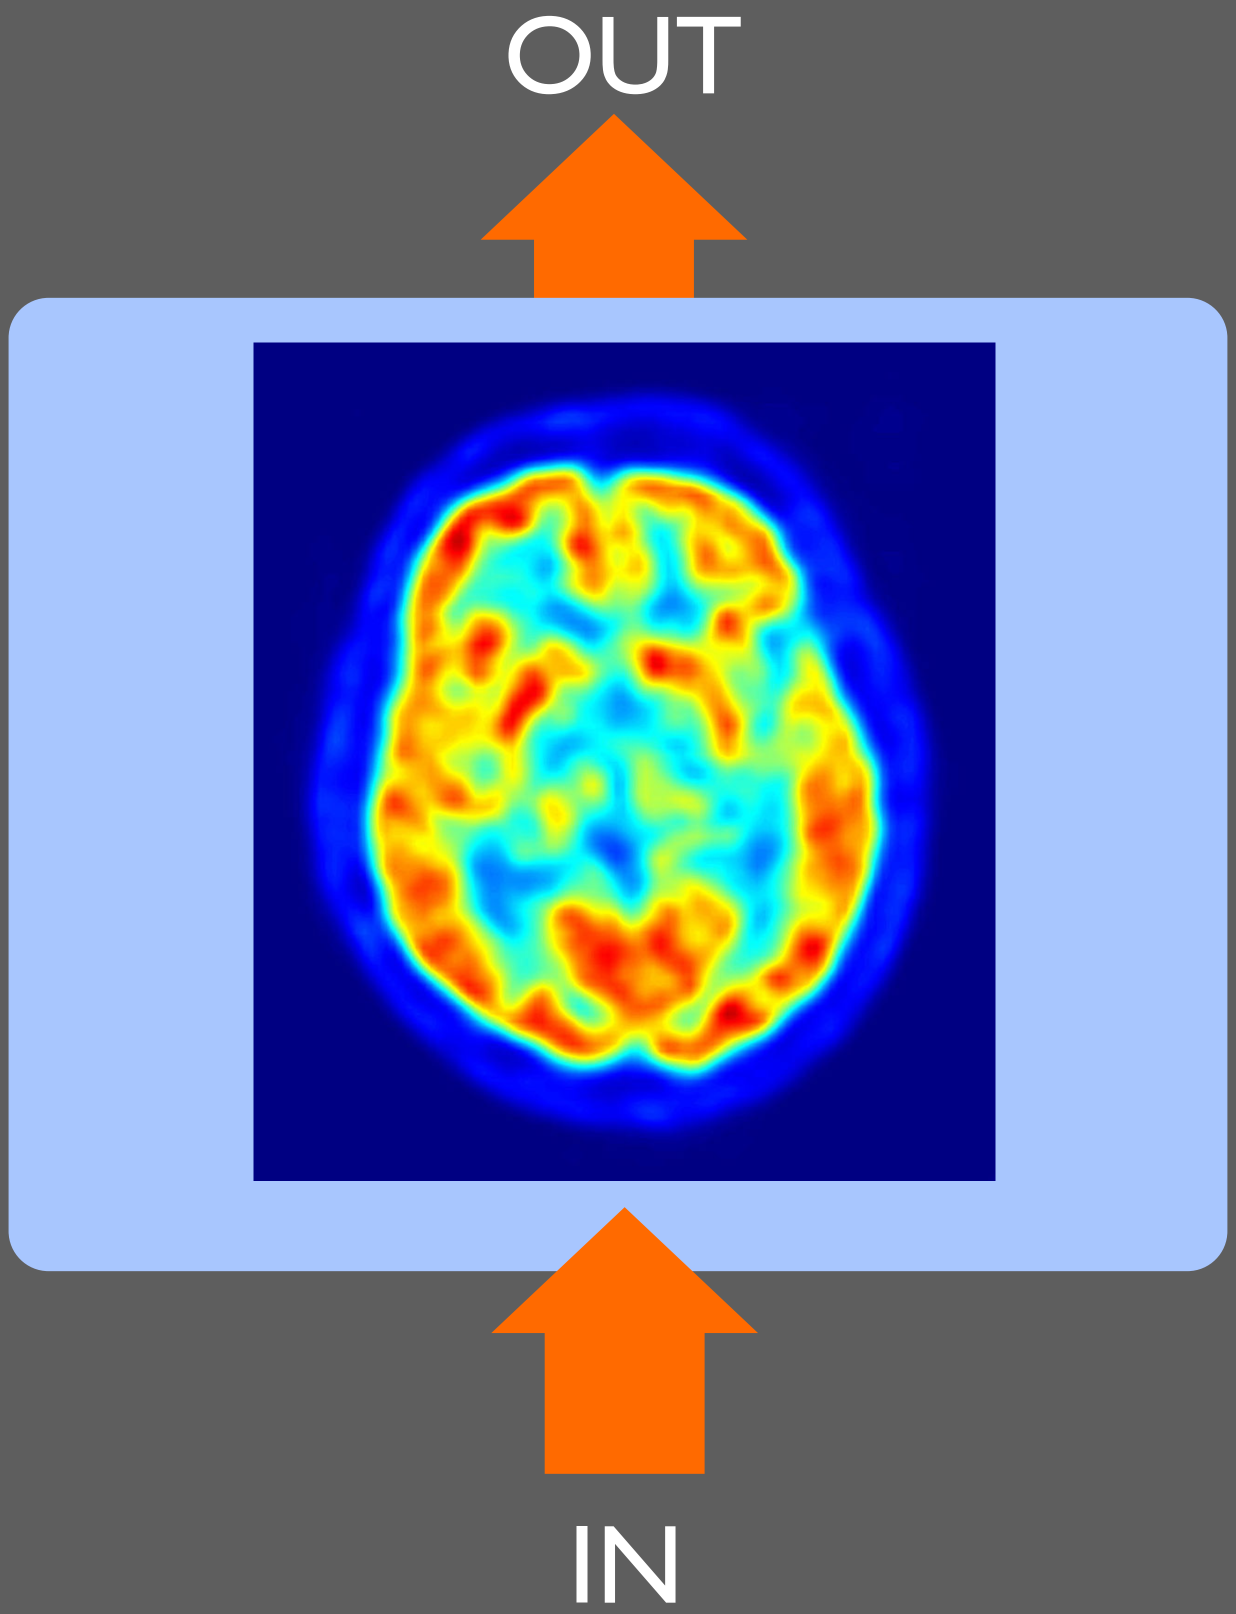
OUT
IN

“light bulb”

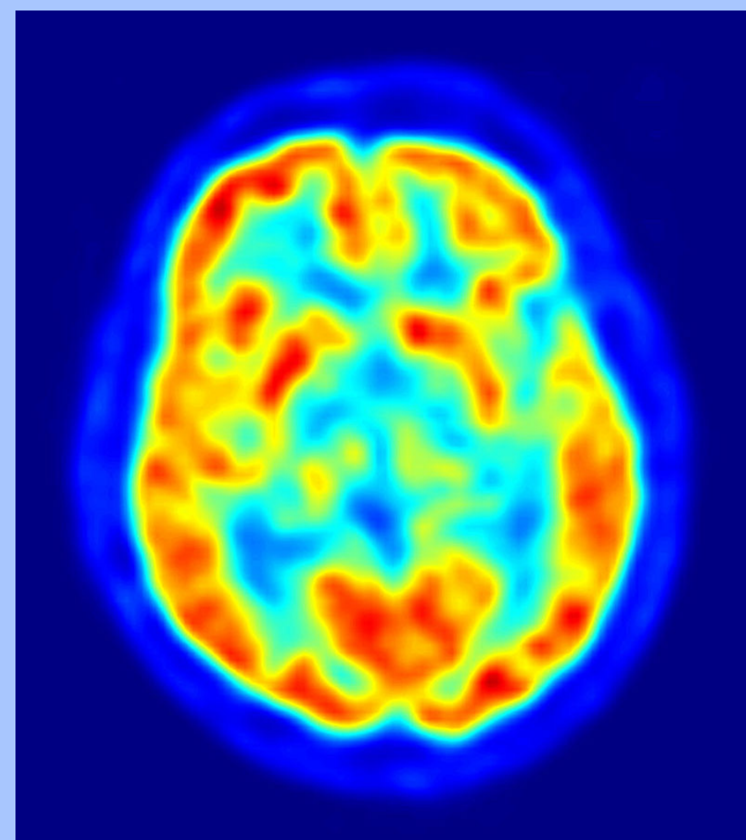




(Pictures: Wikimedia Commons)

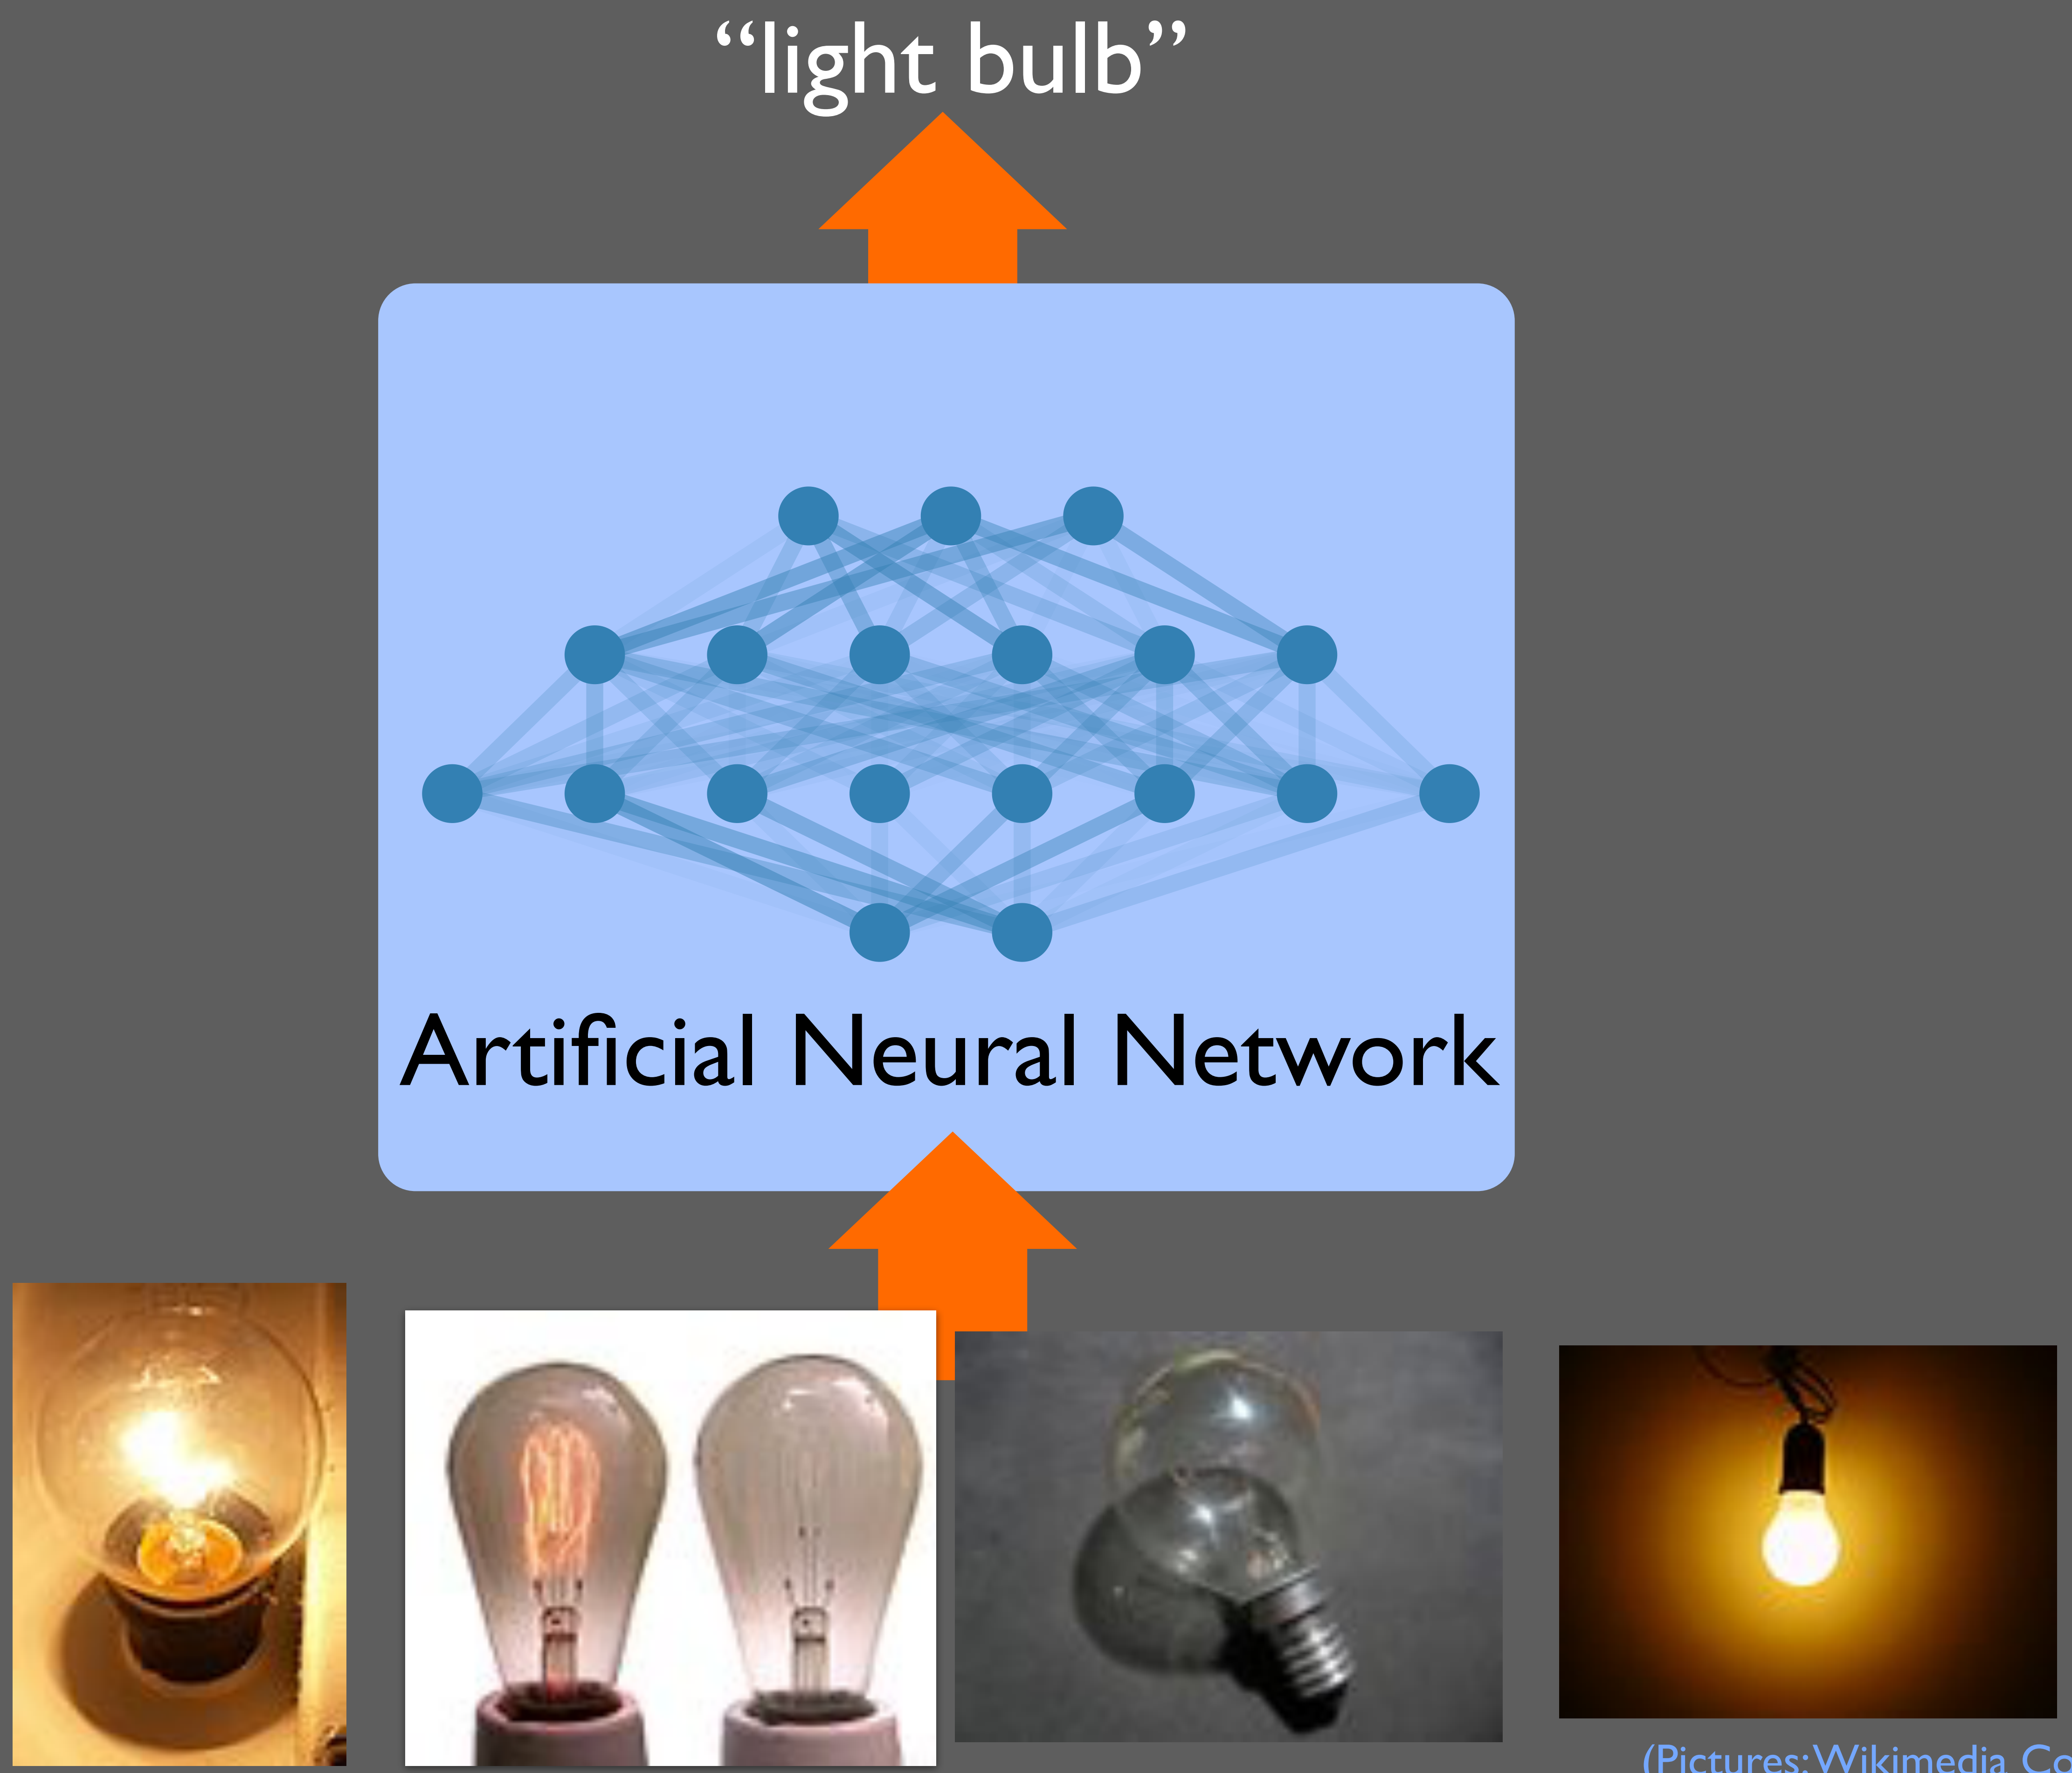

(Pictures: Wikimedia Commons)

2012: A deep neural network beats other approaches clearly, in the “ImageNet” competition

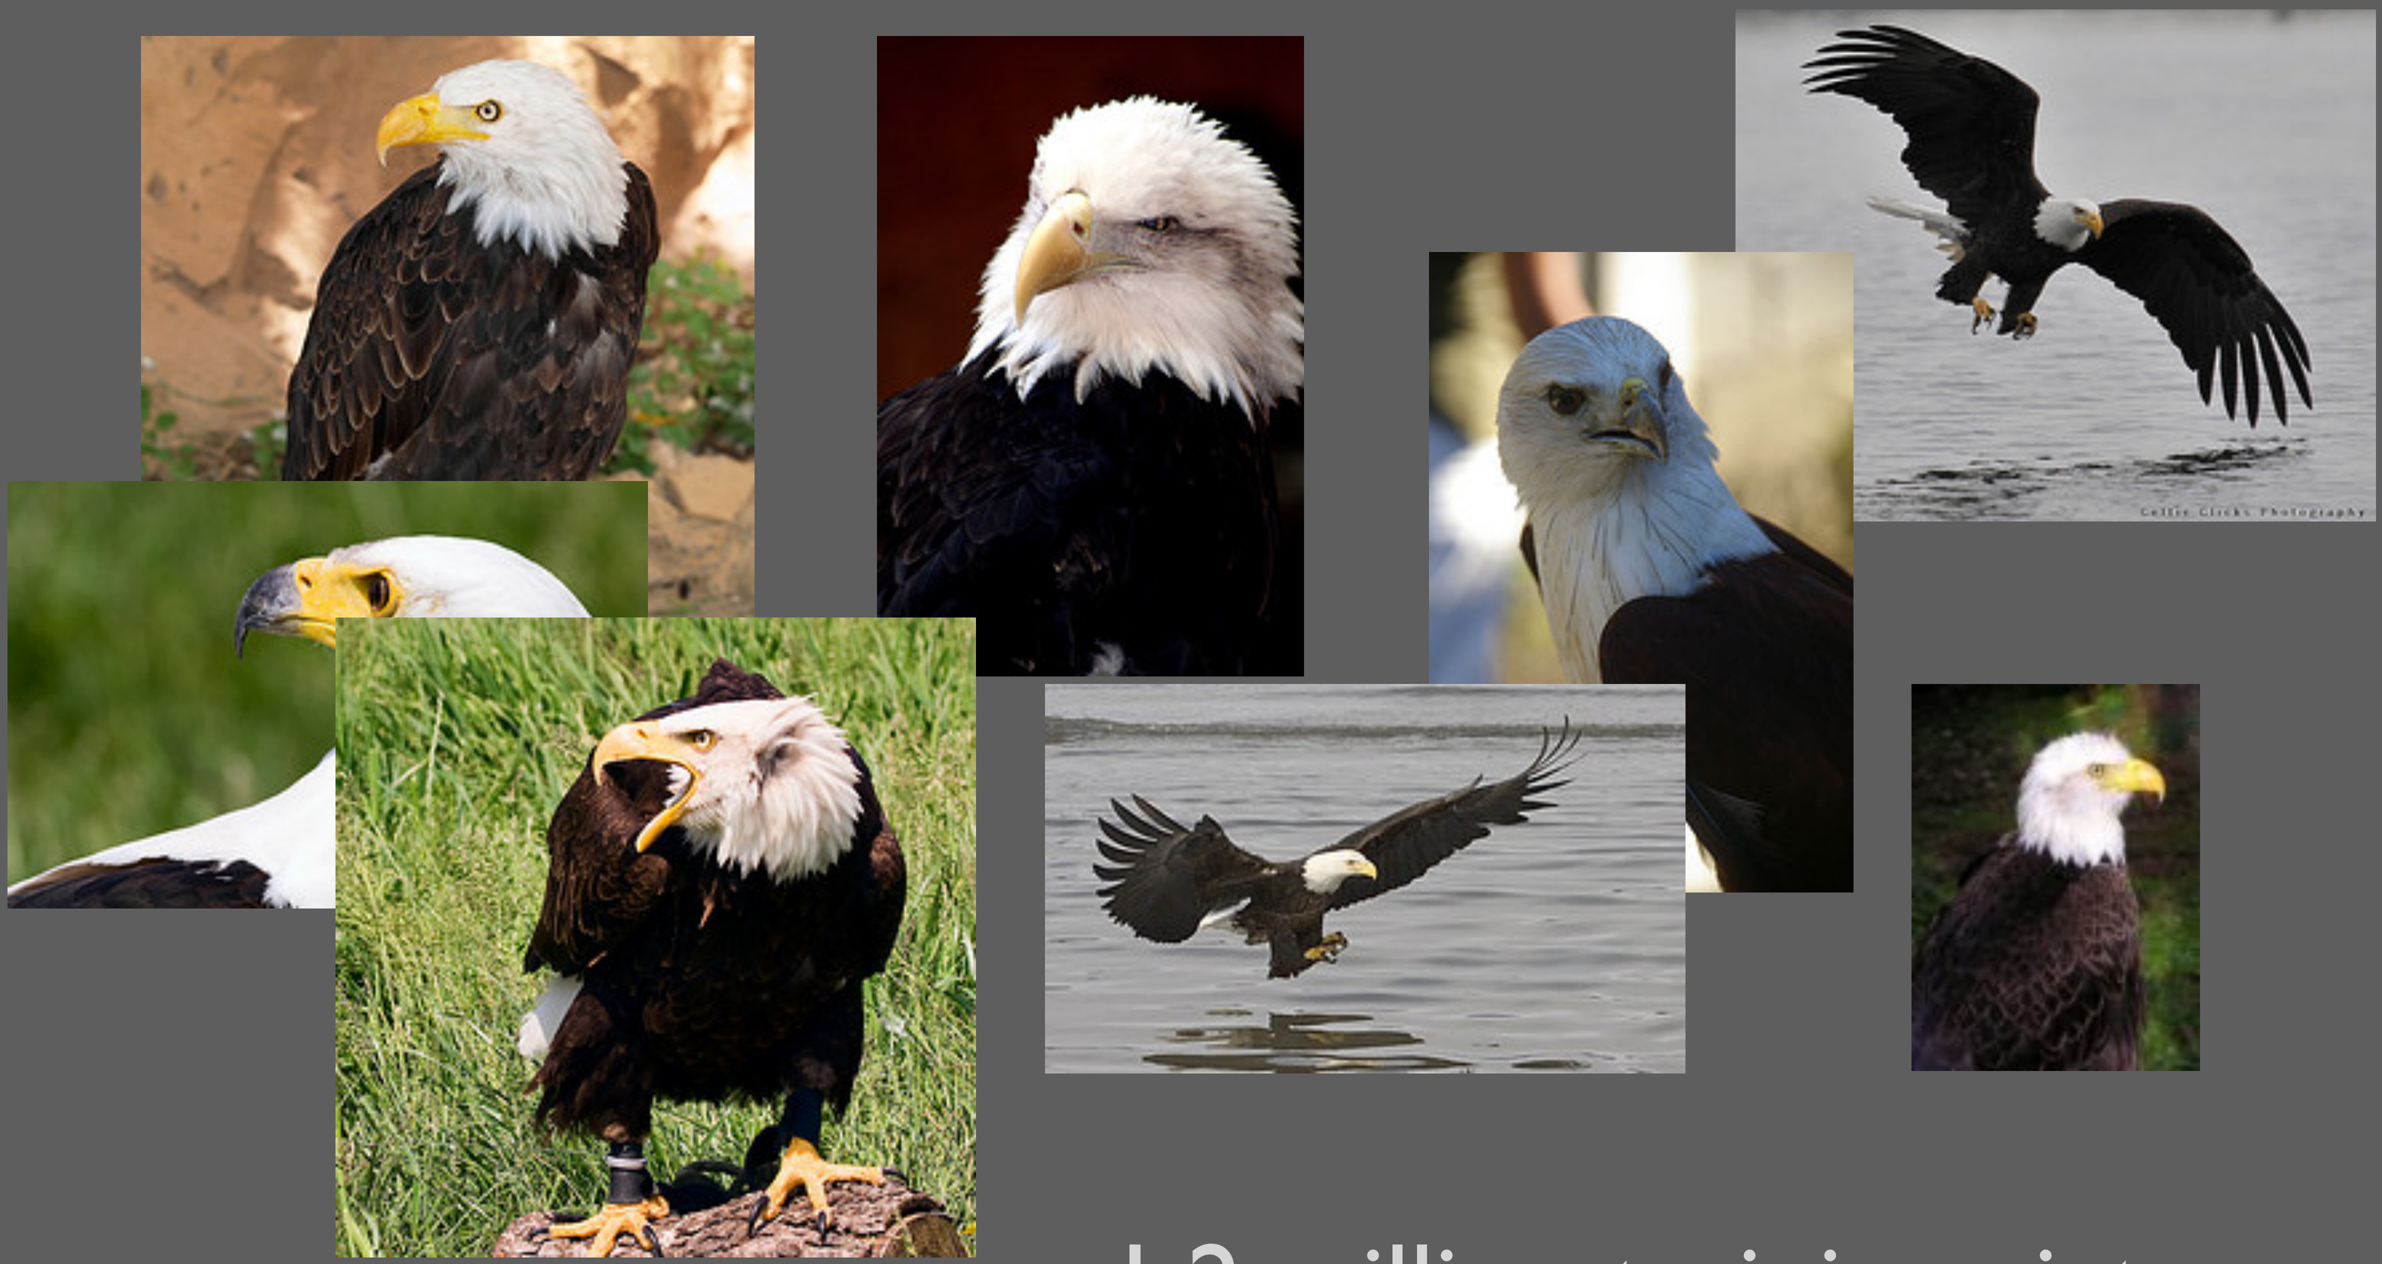

Pictures: image-net.org

1.2 million training pictures (annotated by humans) since 2012: rapid proliferation of real-world applications of artificial neural networks

image labeling, translation, speech recognition, ...

since ~2016: more and more examples in physics

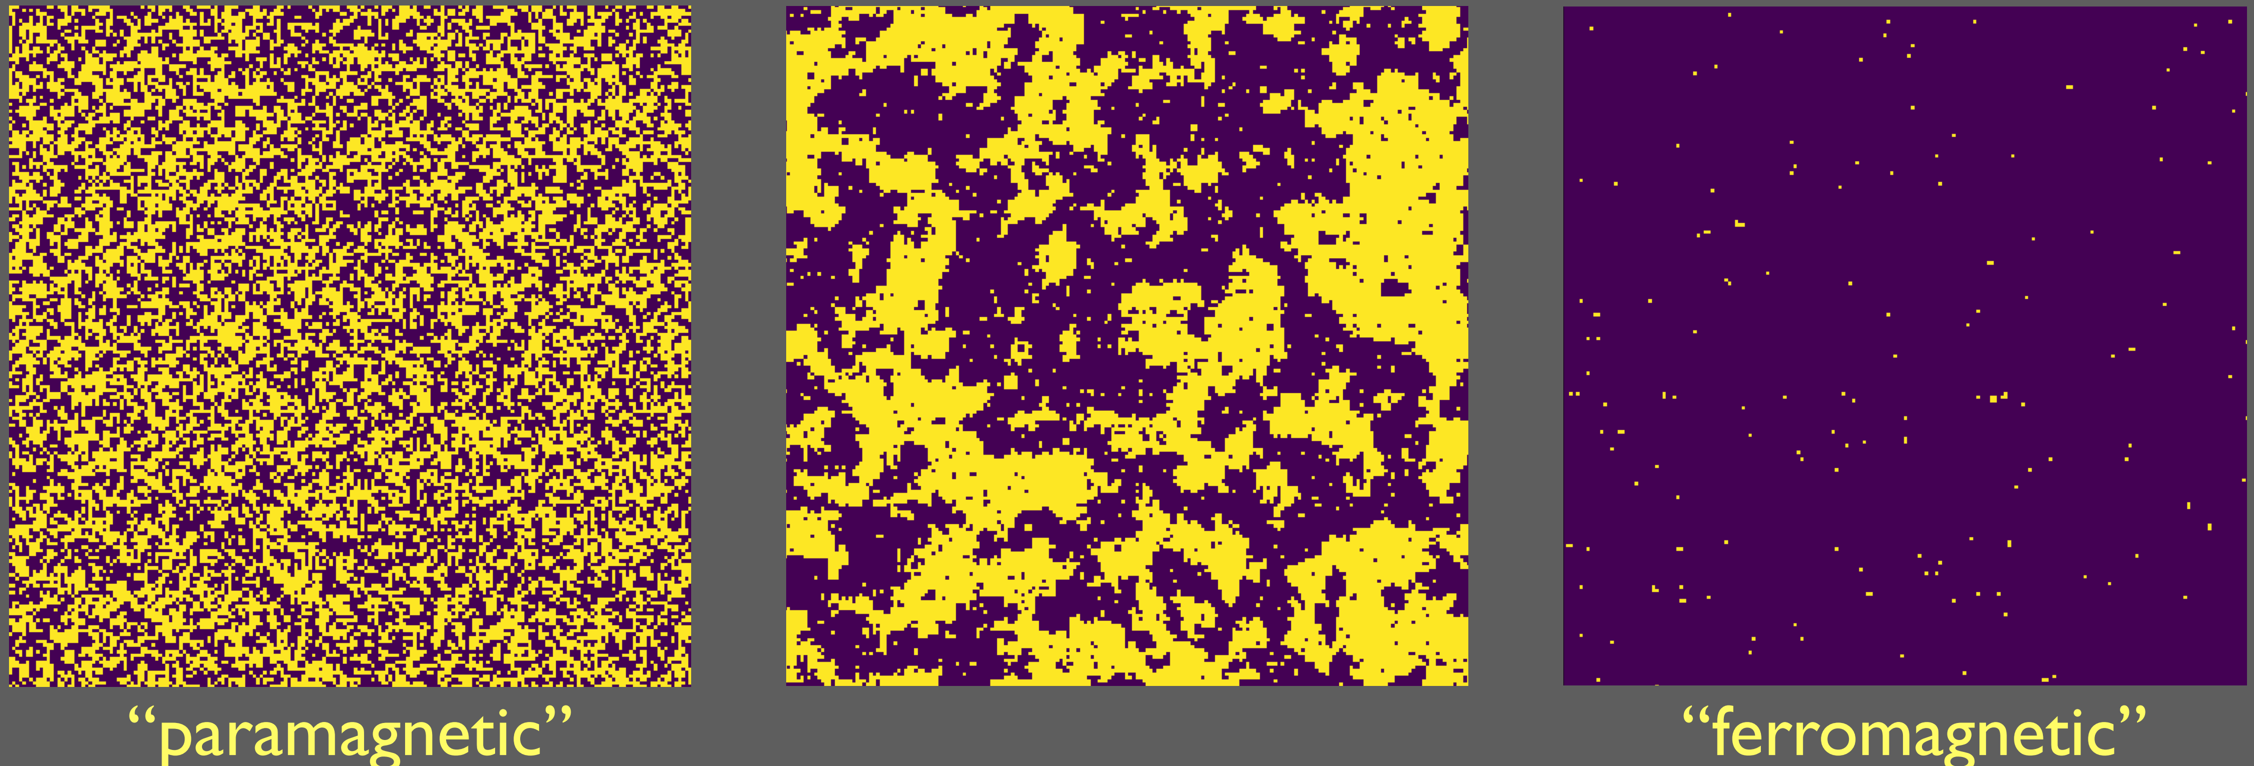

statistical physics, quantum many-body physics, dynamical systems, experimental data analysis, …

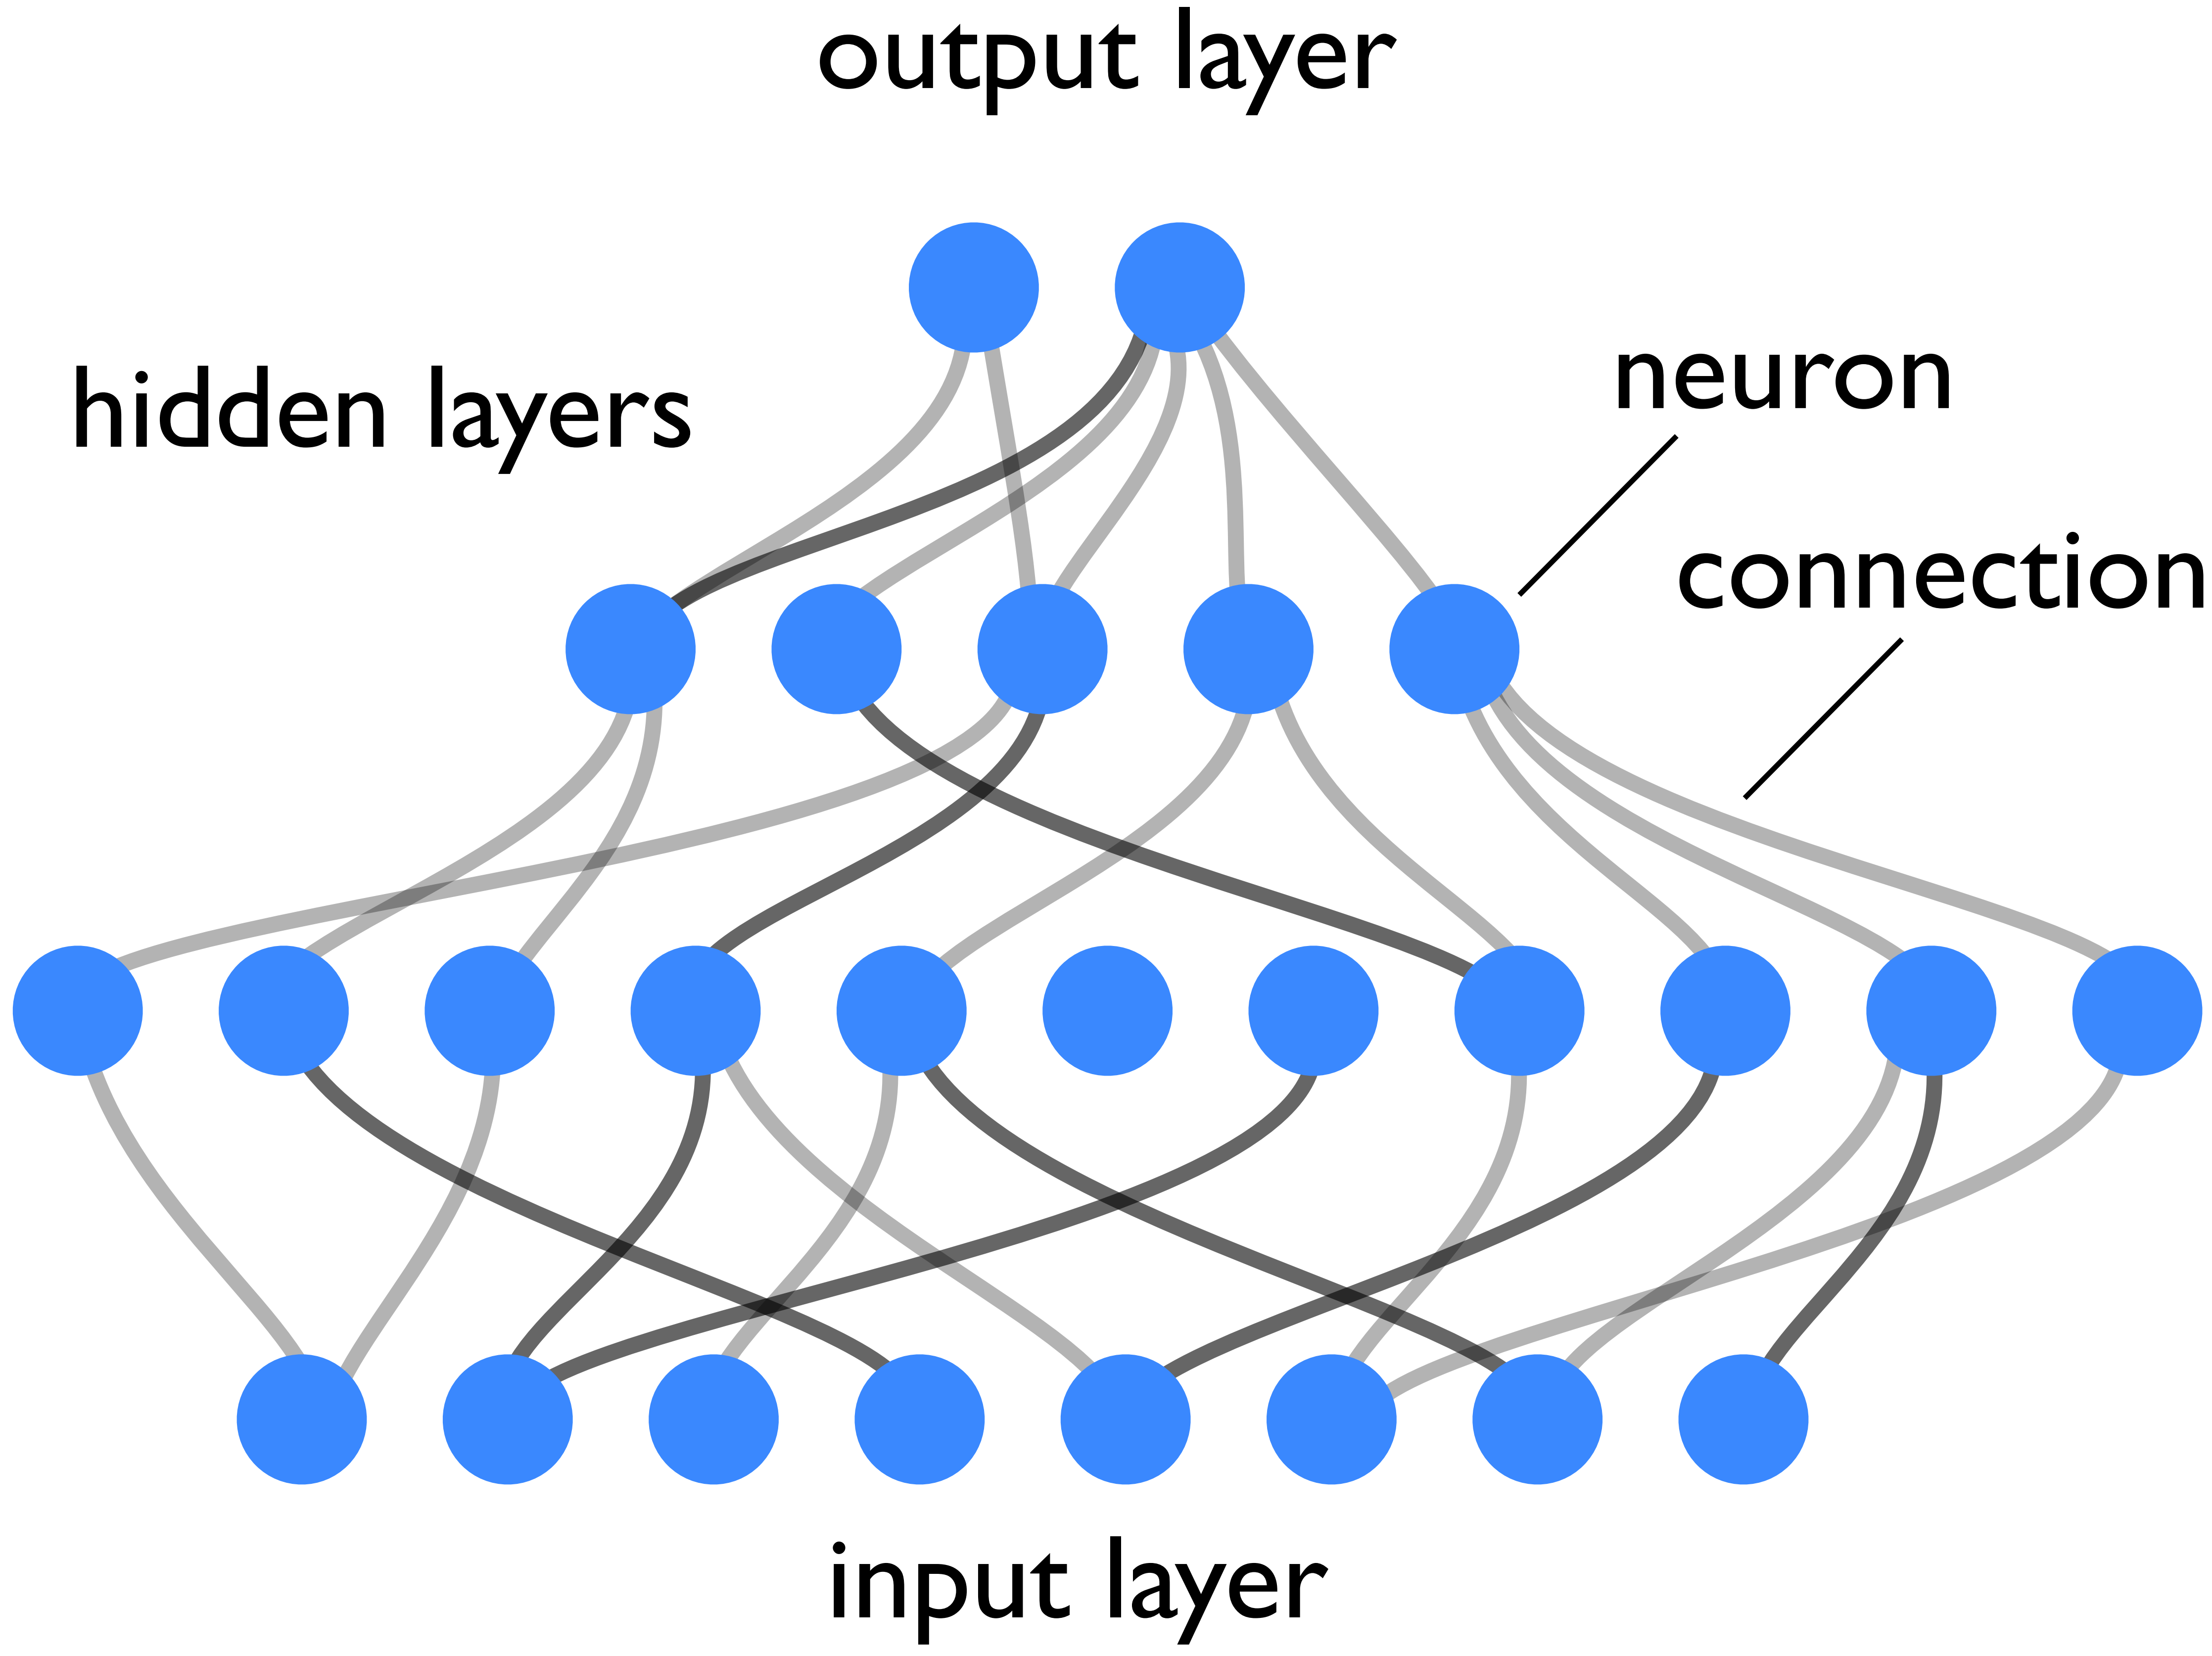
output layer
hidden layers
neuron
connection
input layer

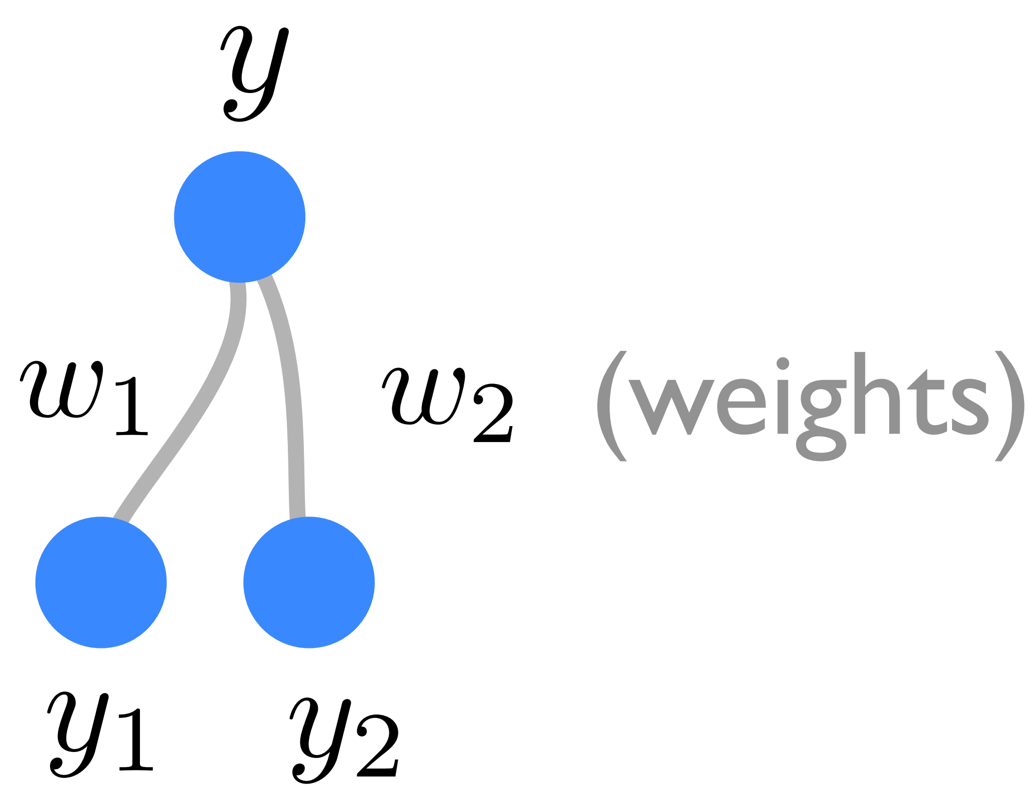

1 linear superposition

$$z = w_1 y_1 + w_2 y_2 + b$$

2 apply nonlinear activation function

$$y = f(z)$$

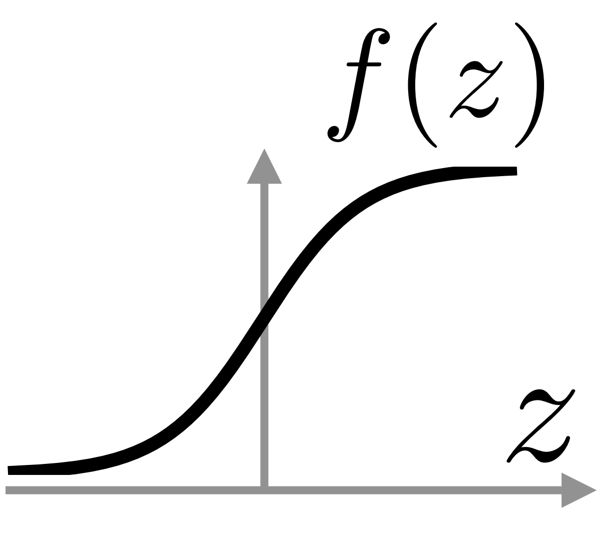

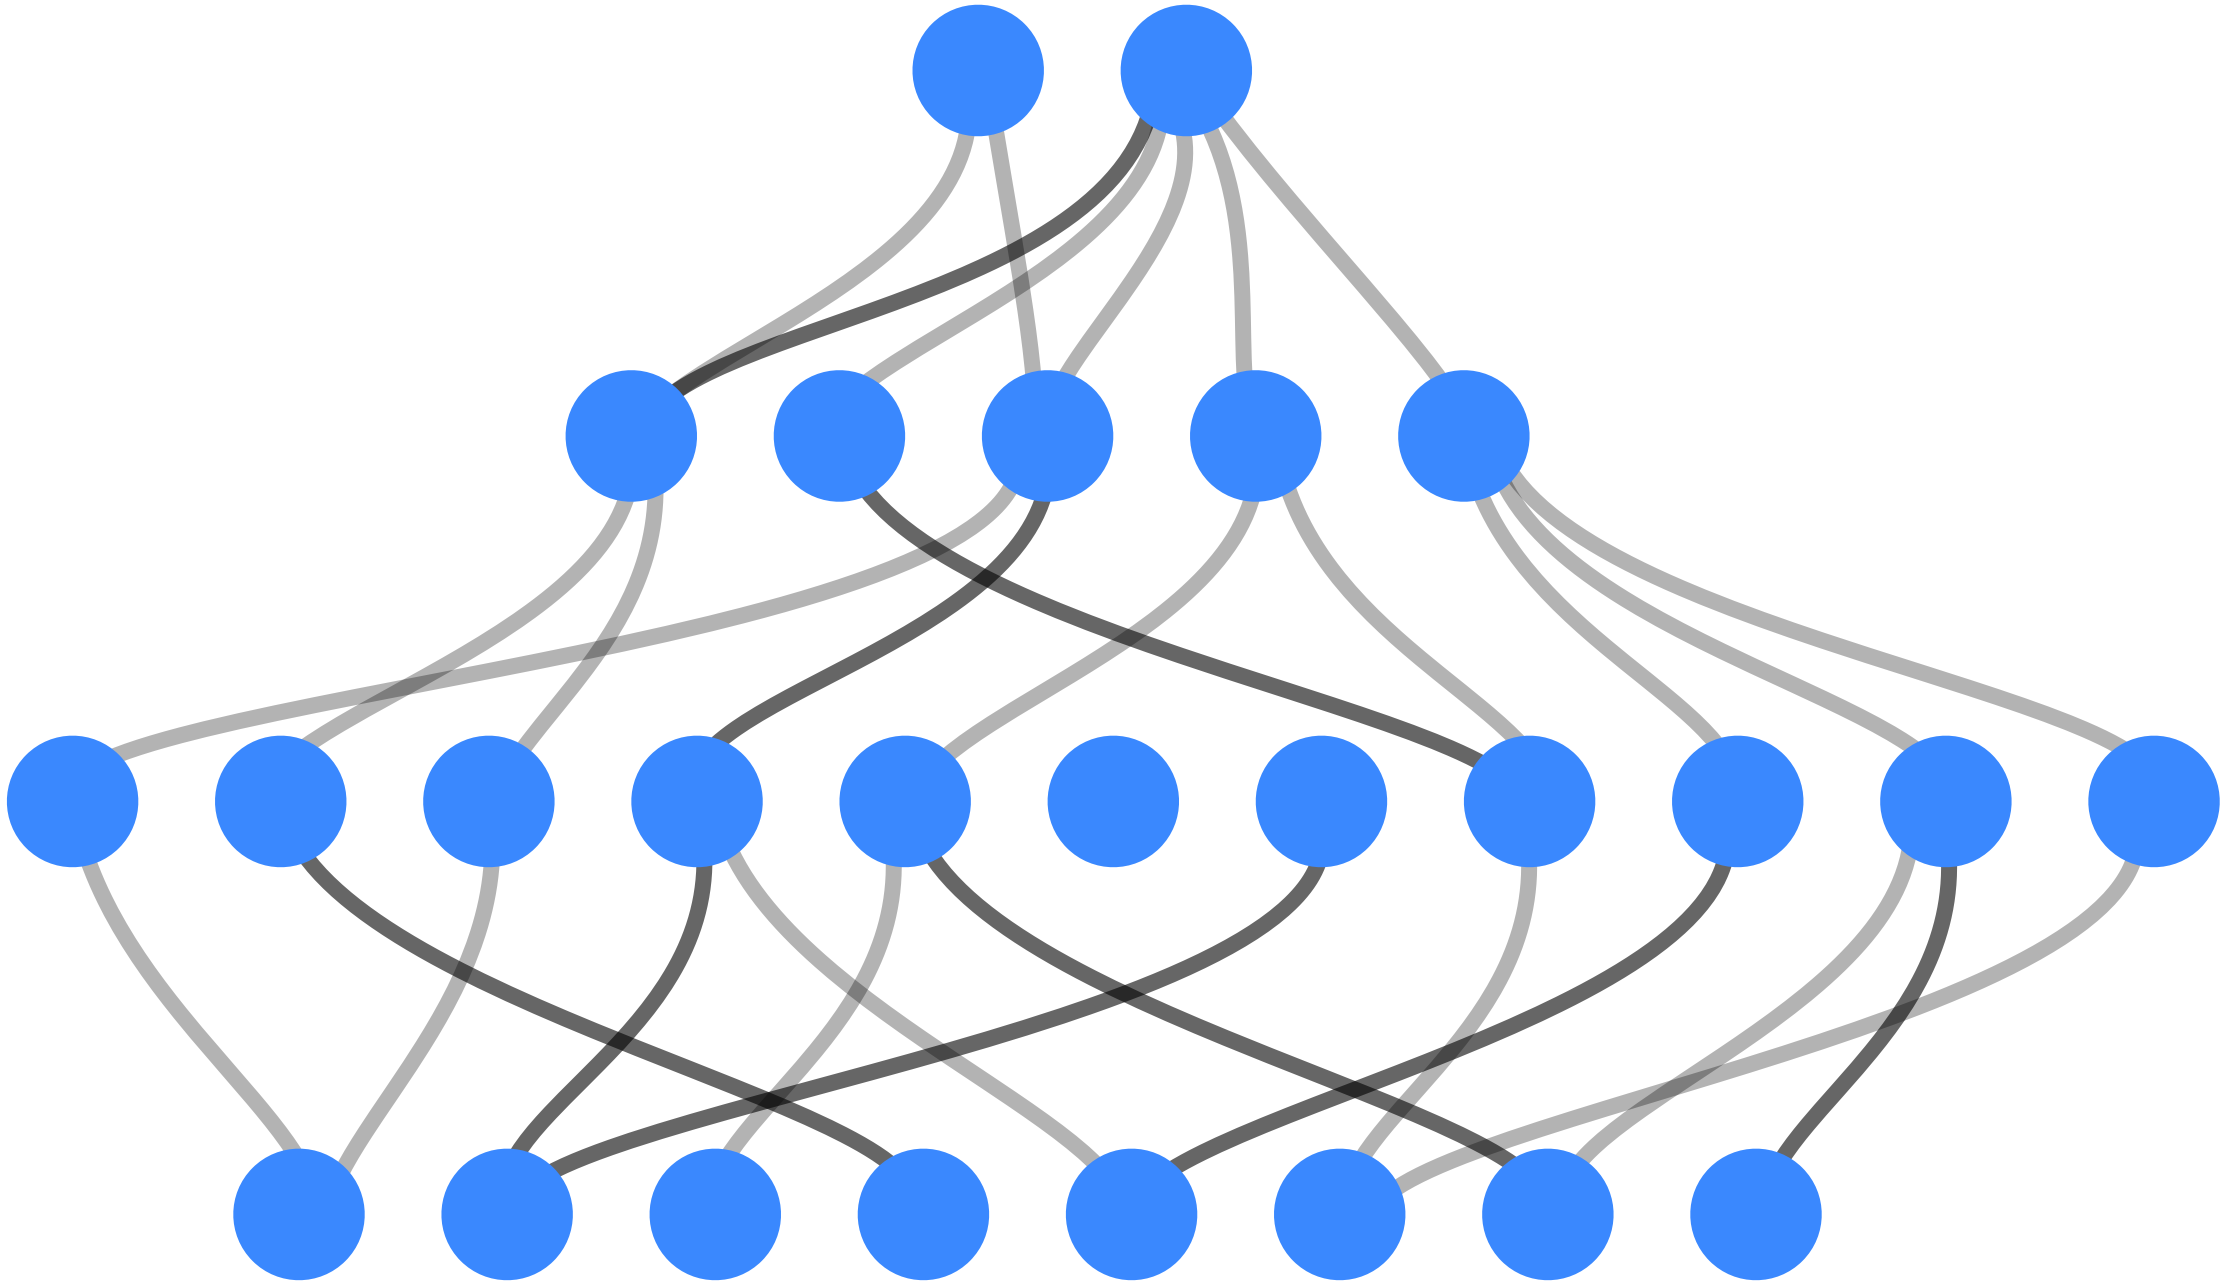

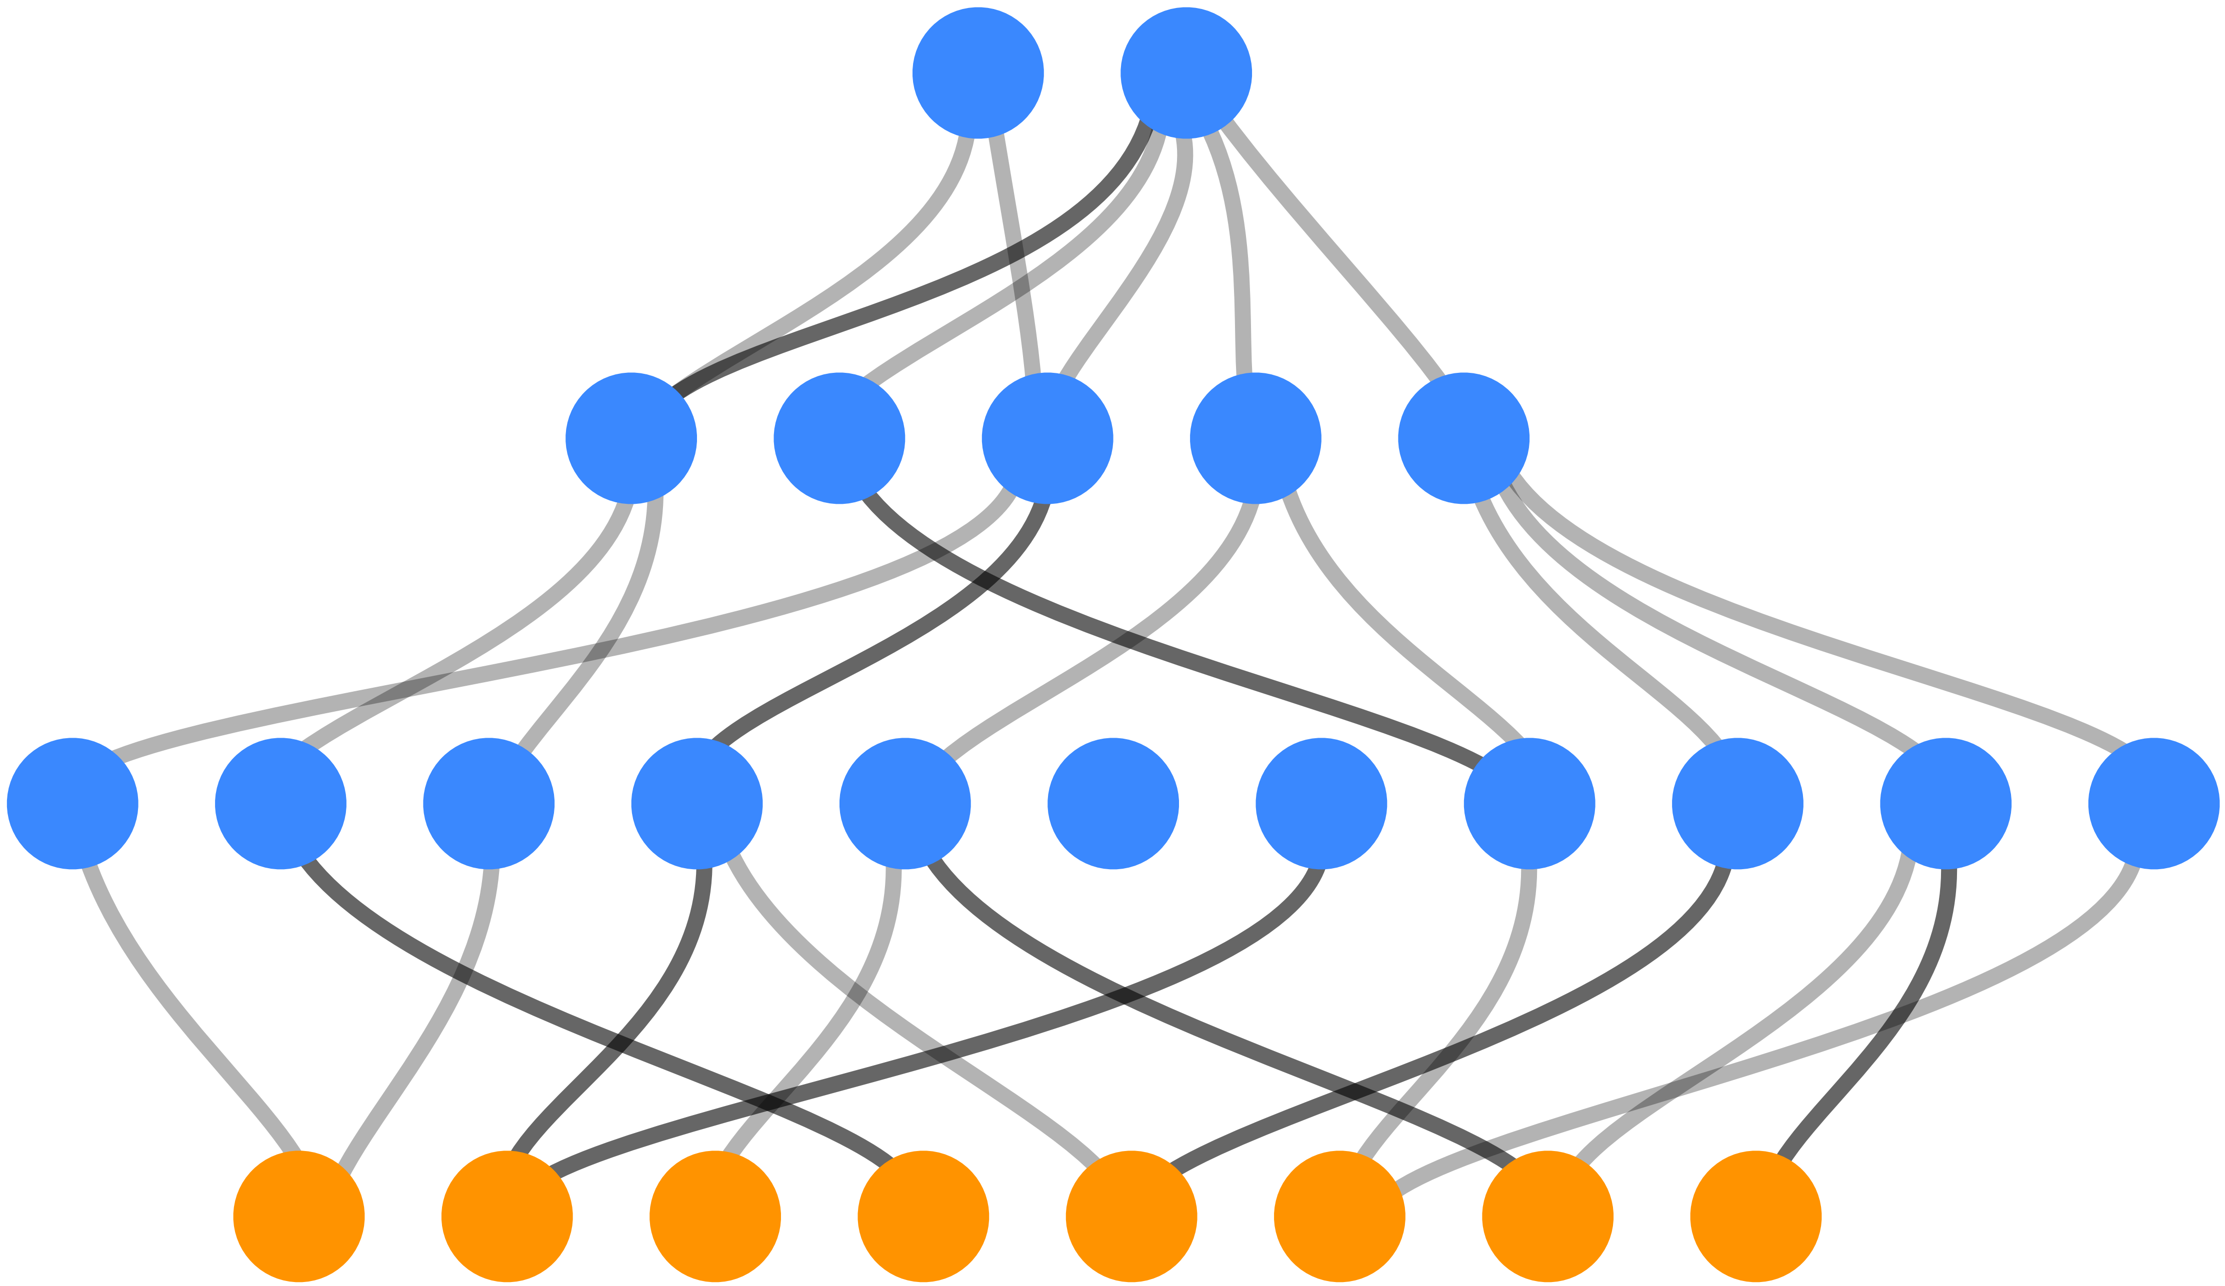

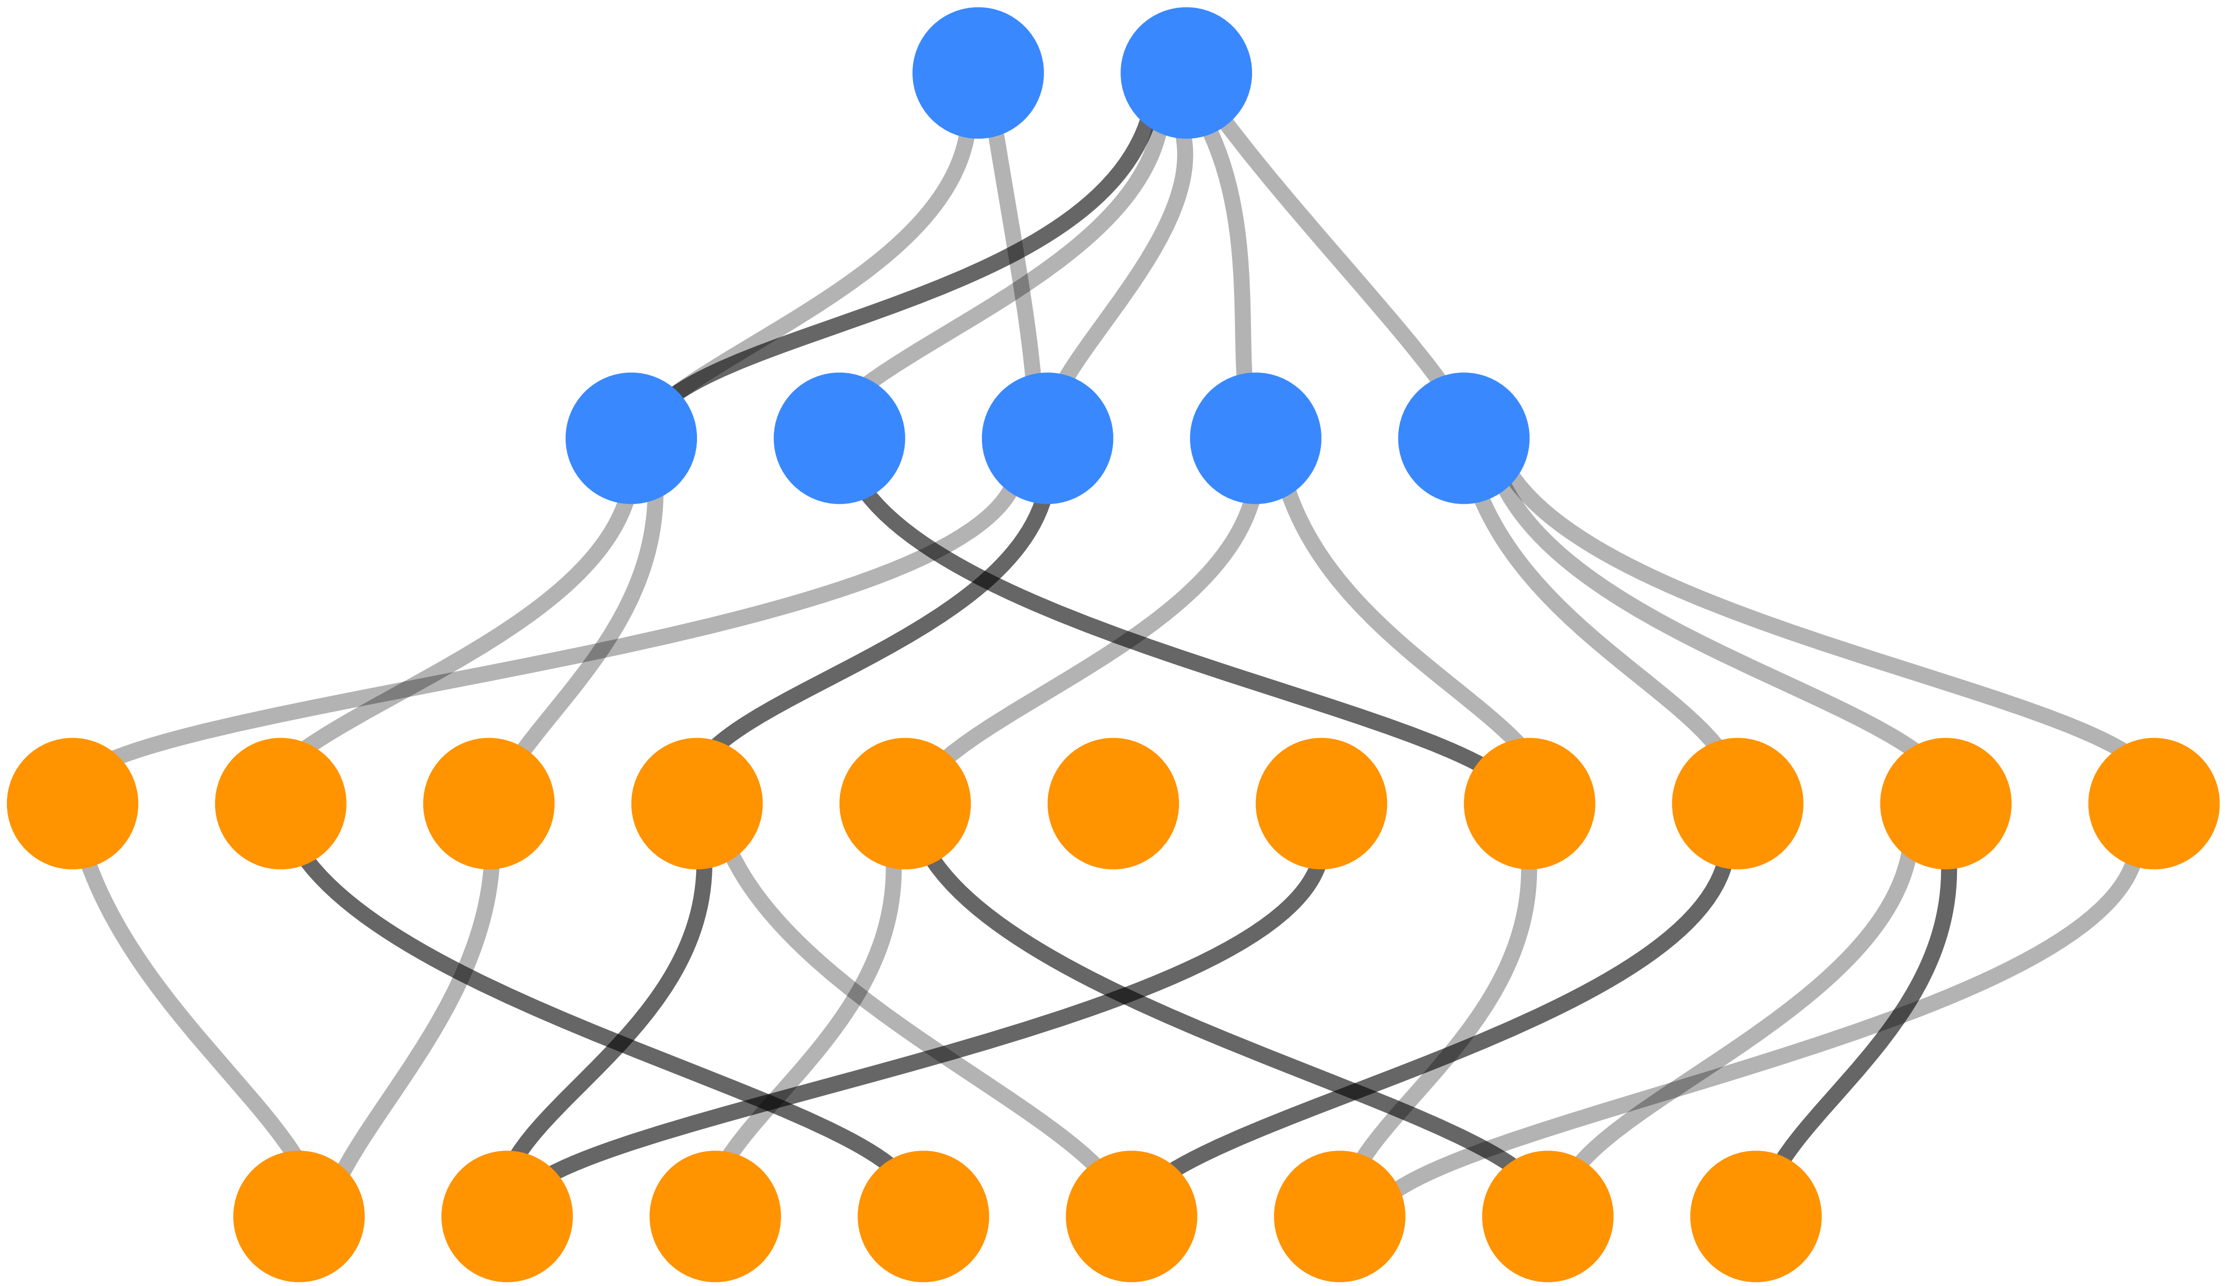

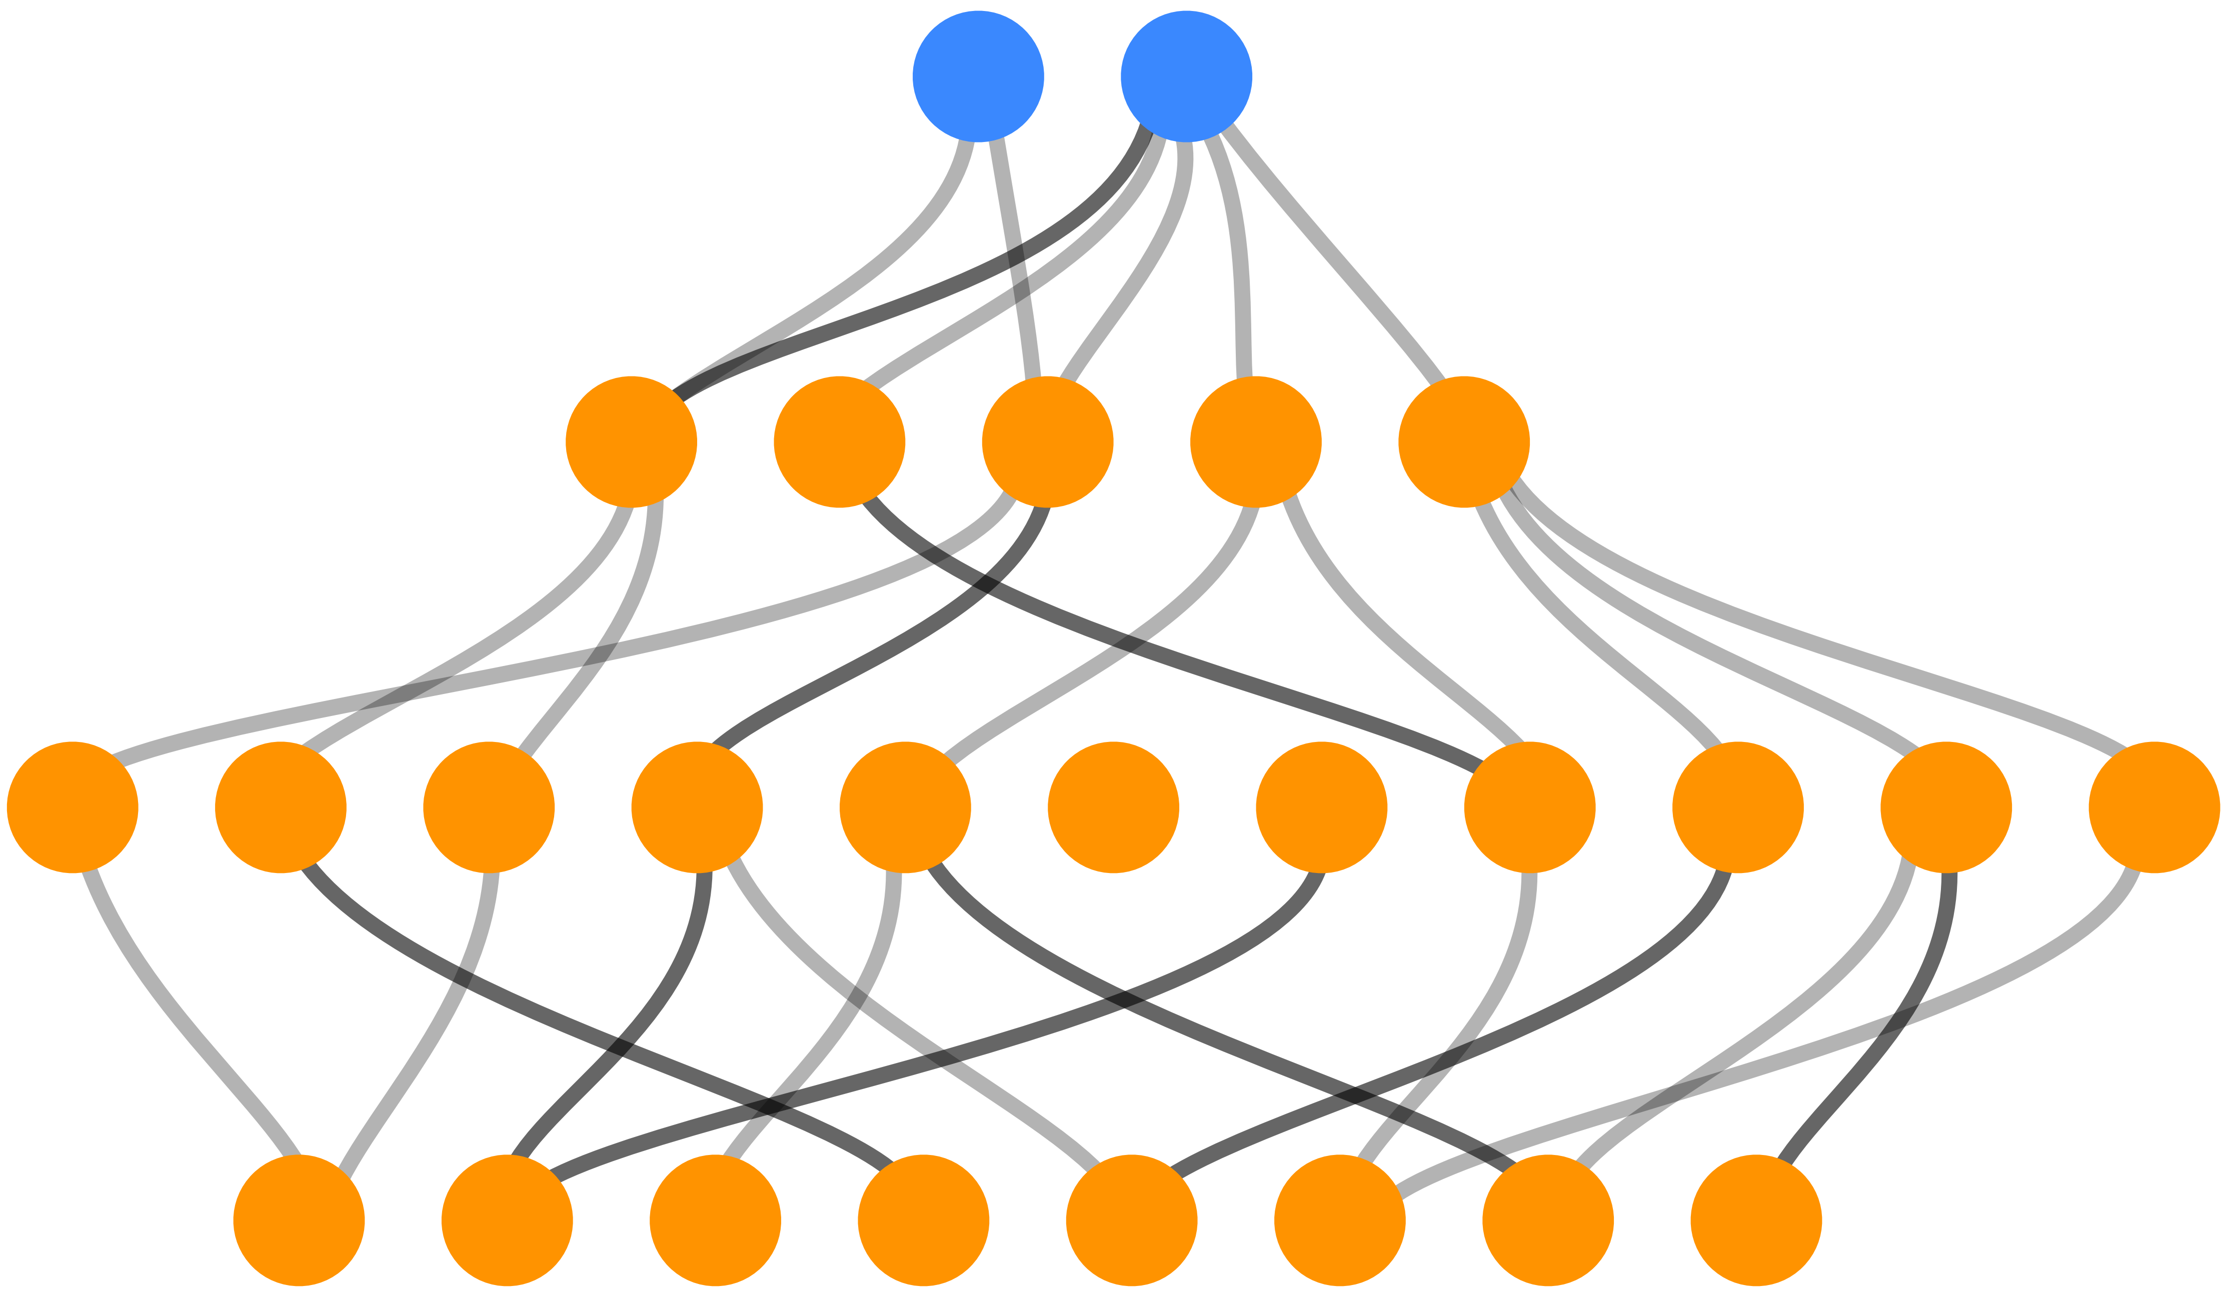

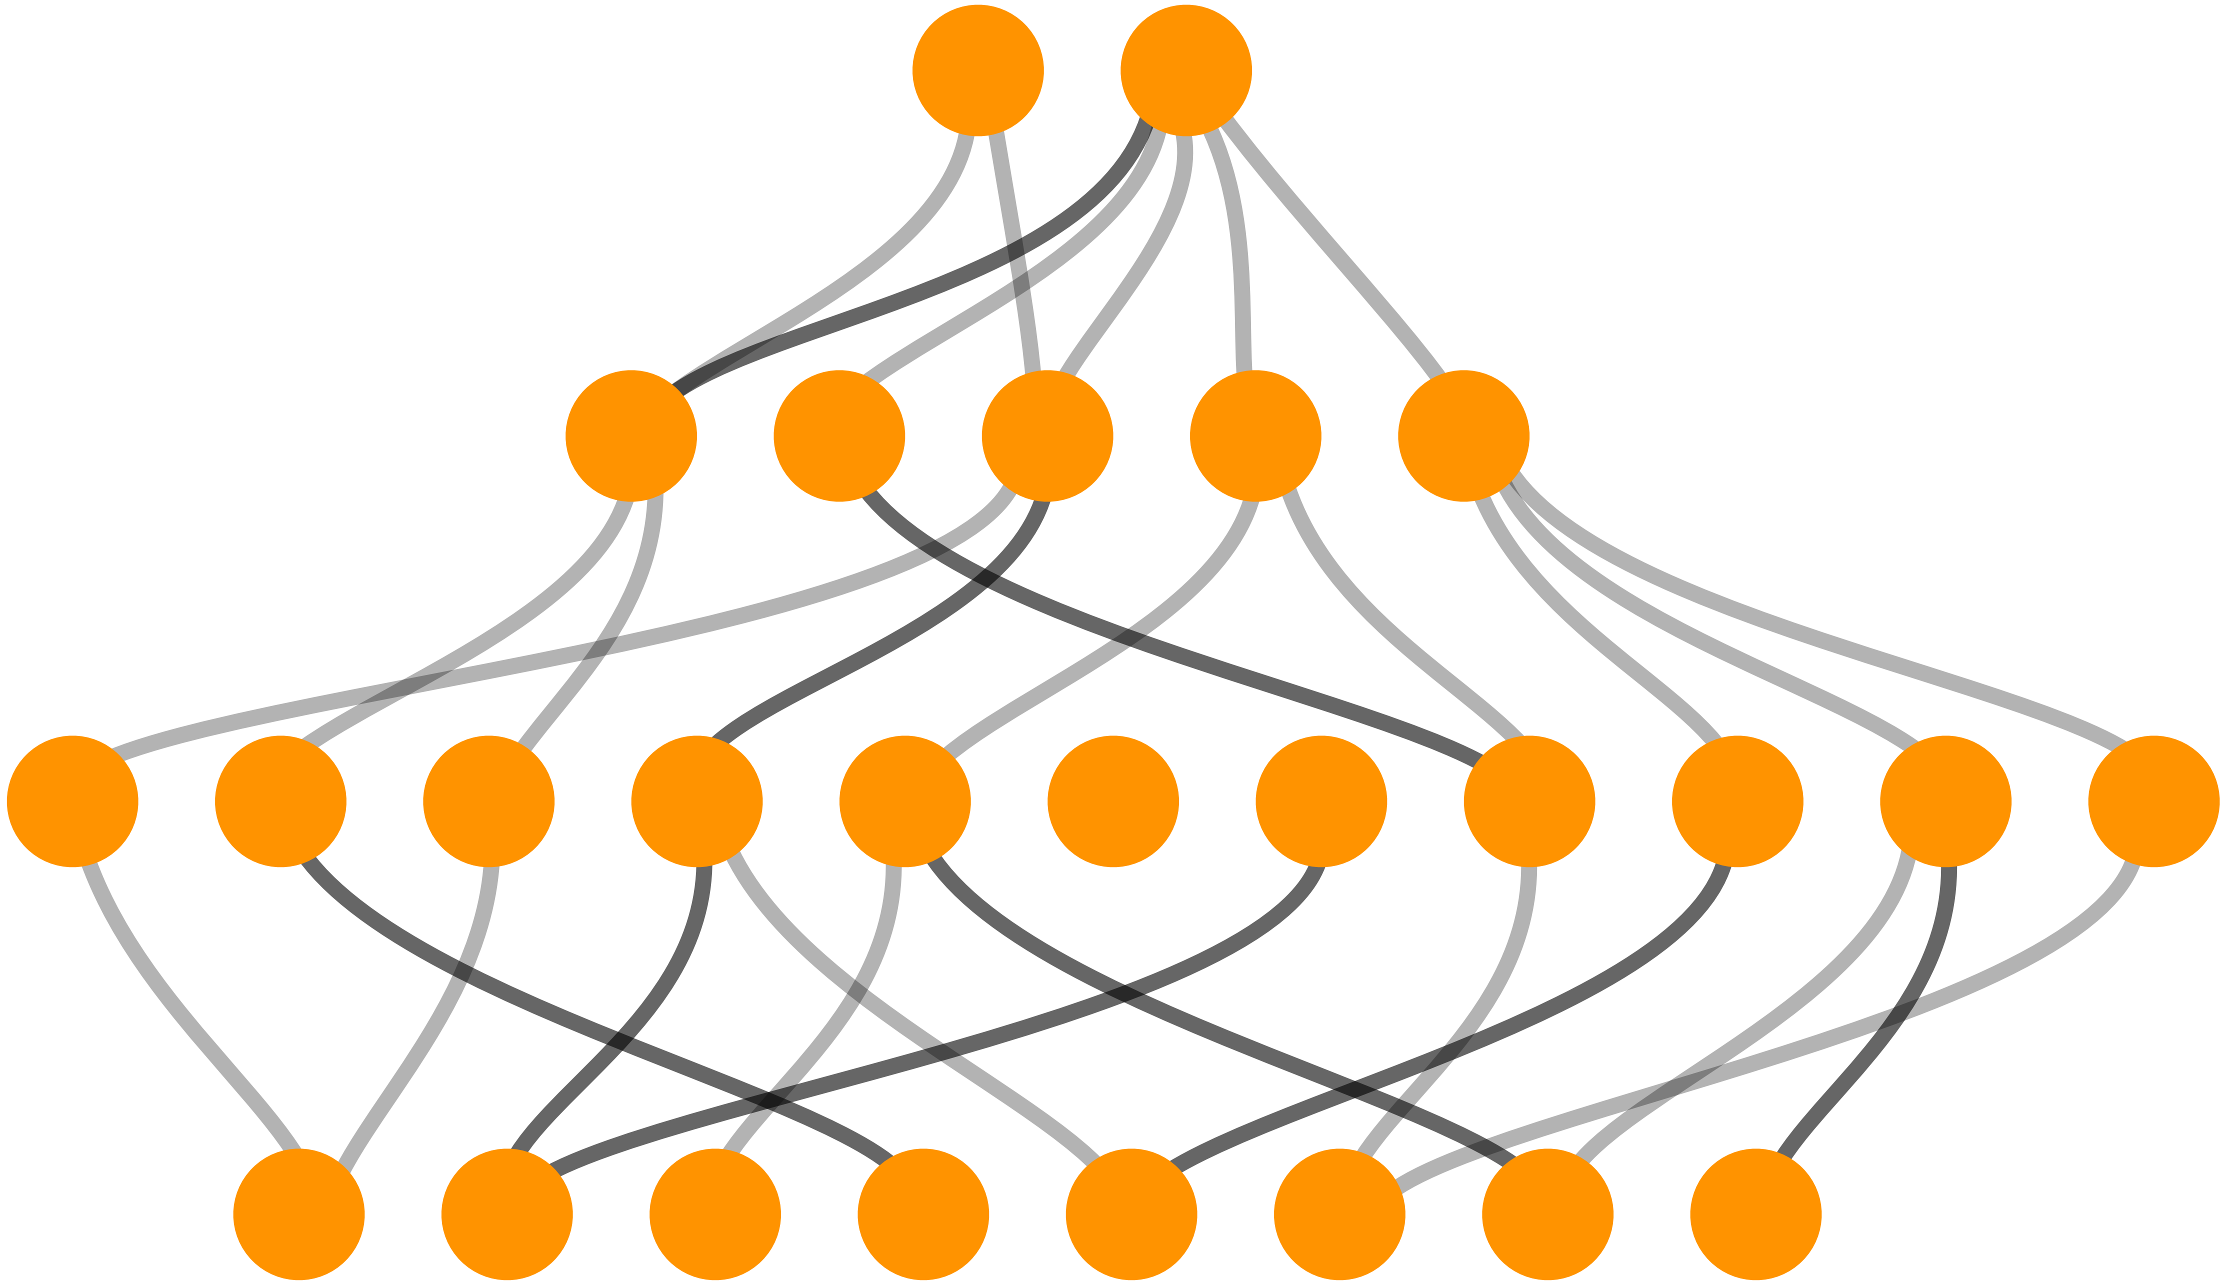

Training: adapt weights to come closer to correct input/output relation, based on thousands of training examples

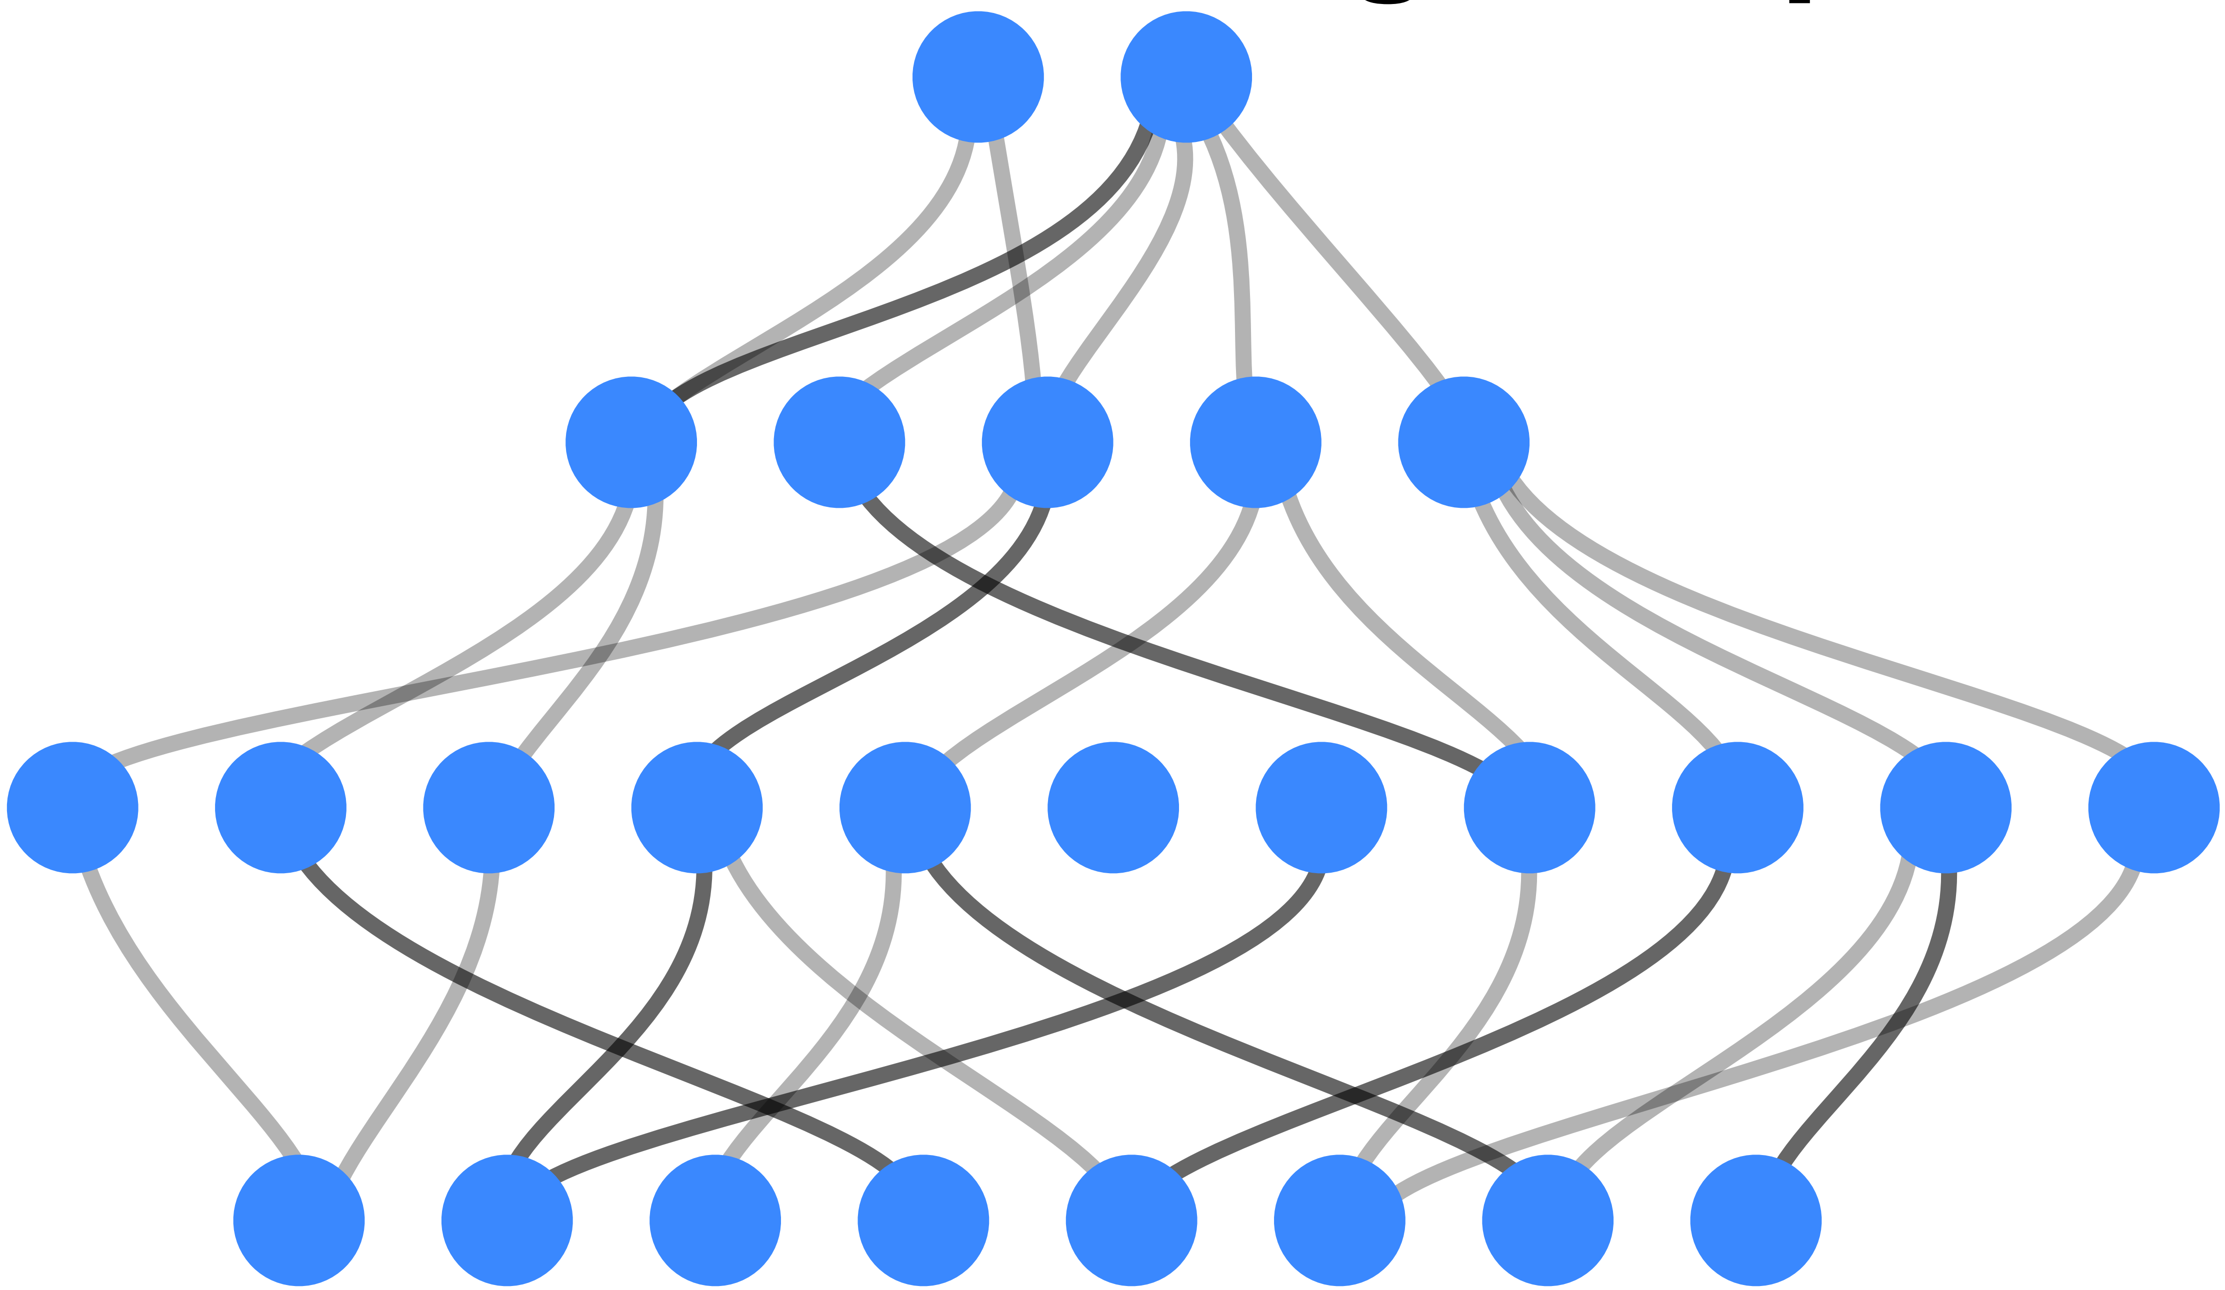

Training: adapt weights to come closer to correct input/output relation, based on thousands of training examples

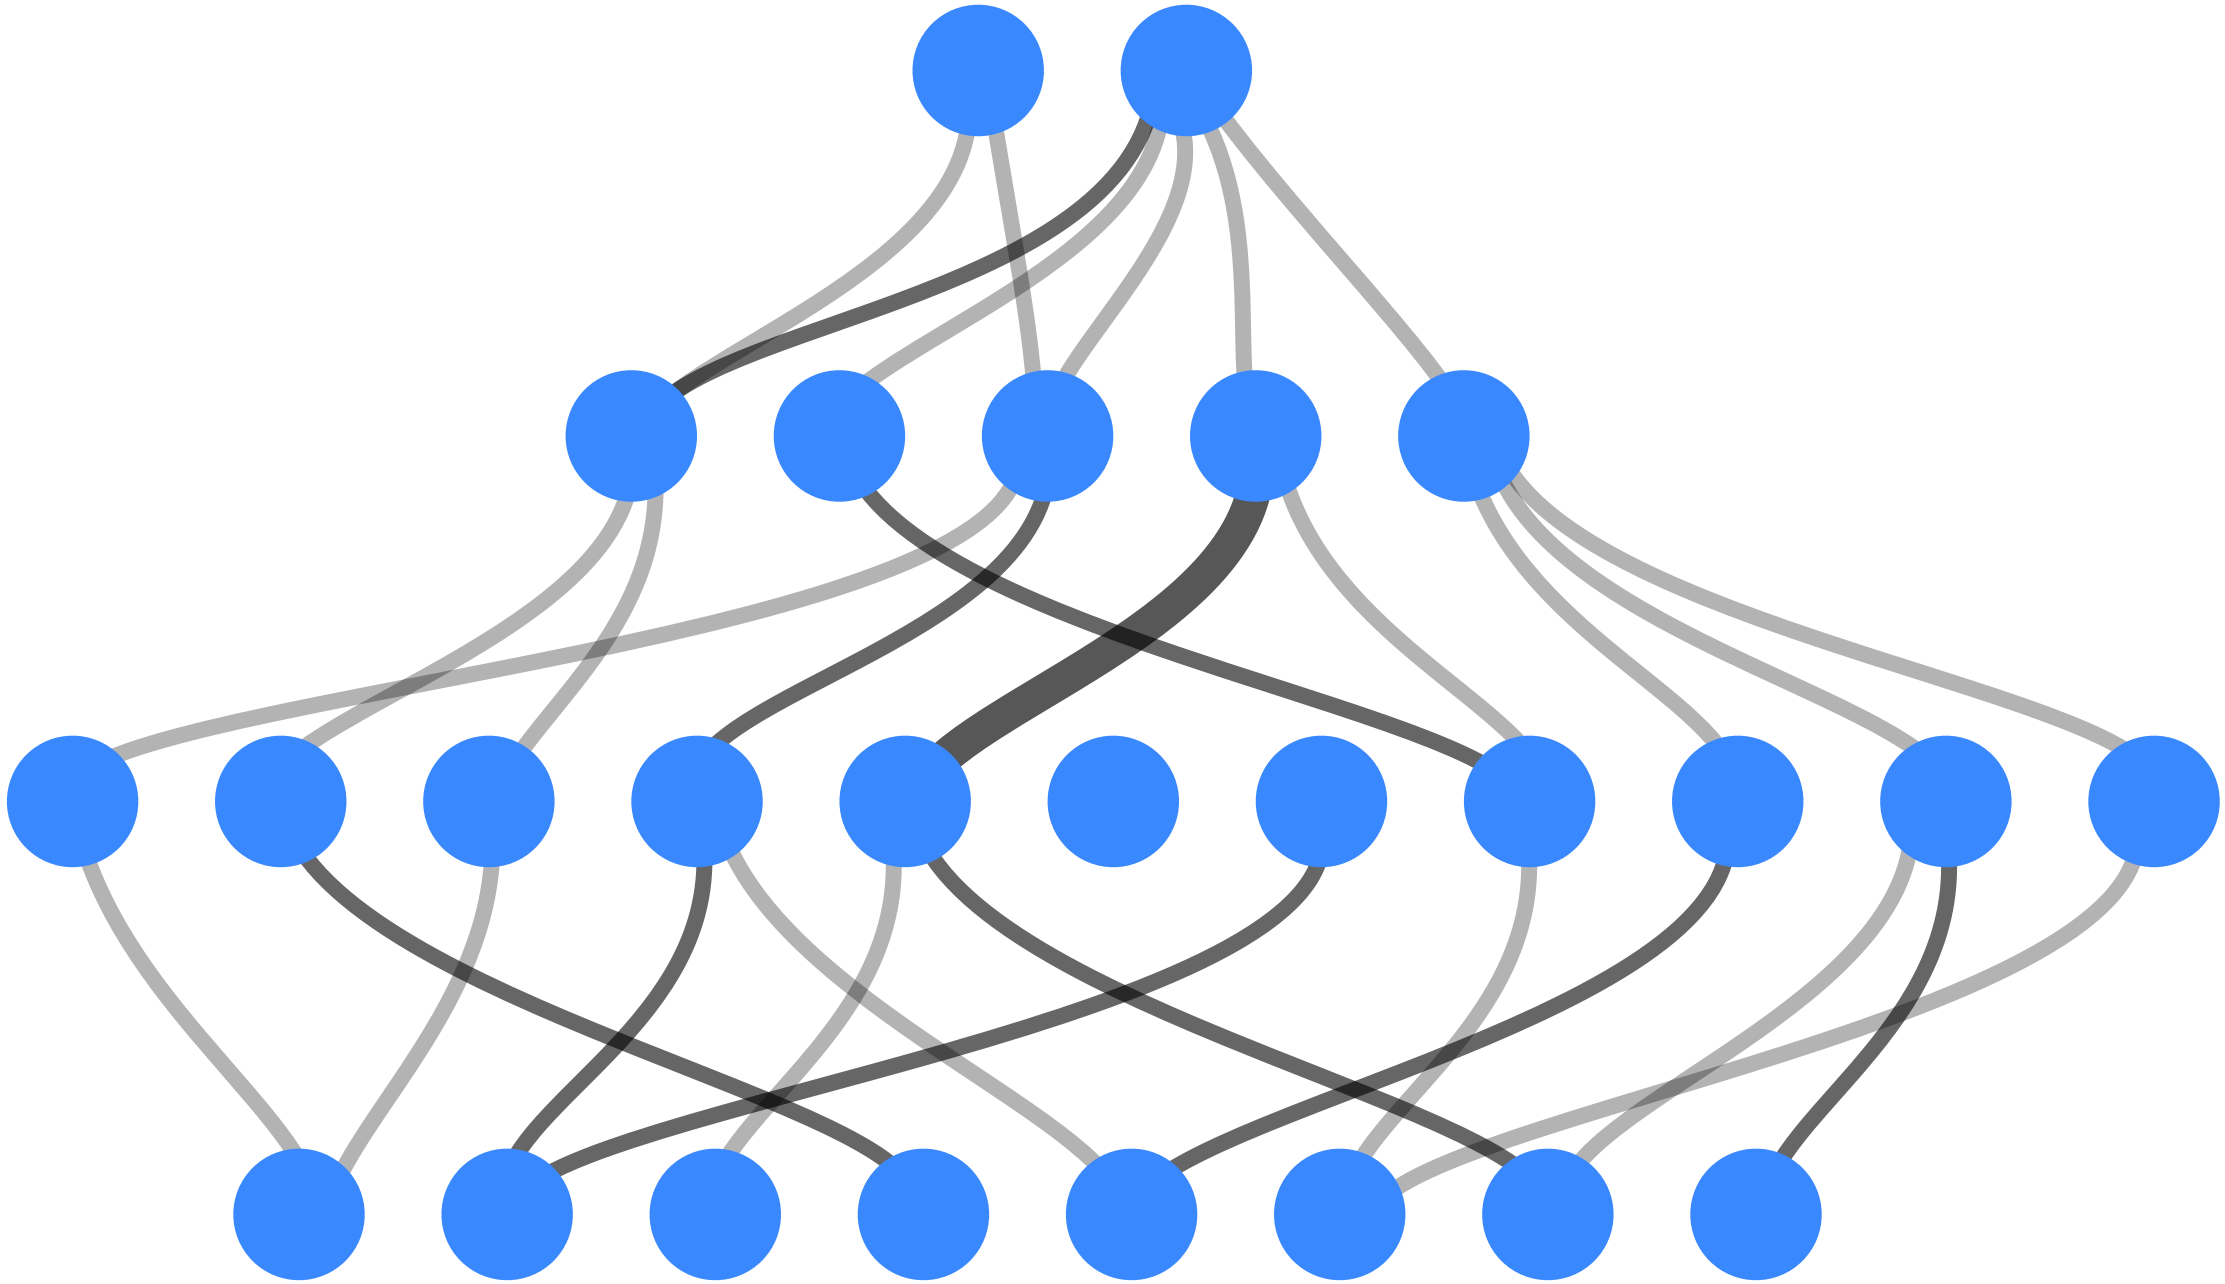

Training: adapt weights to come closer to correct input/output relation, based on thousands of training examples

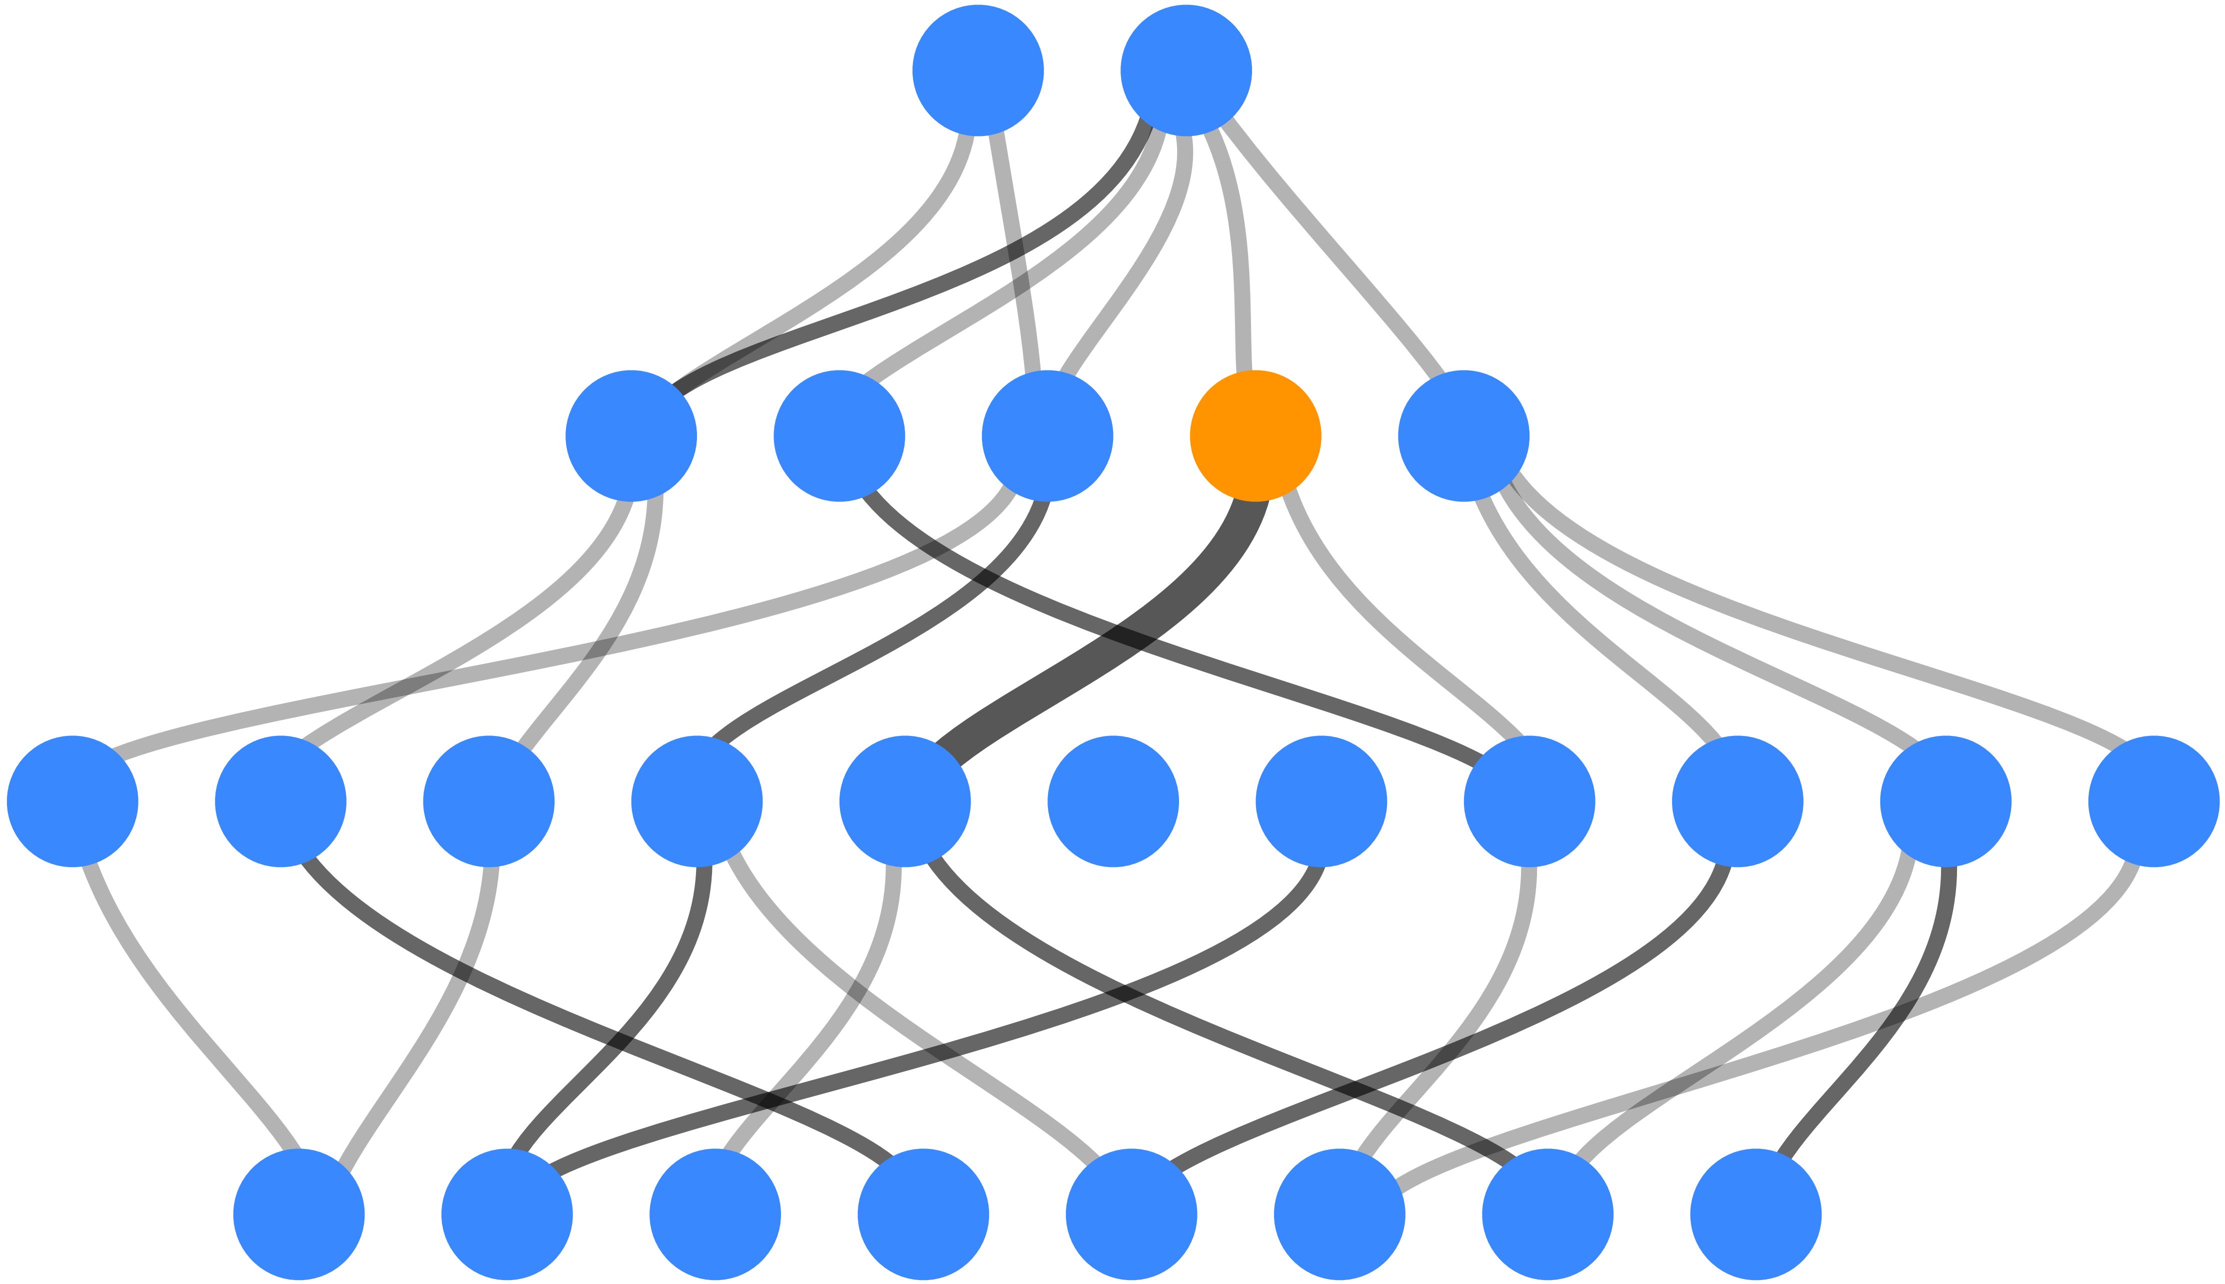

Training: adapt weights to come closer to correct input/output relation, based on thousands of training examples

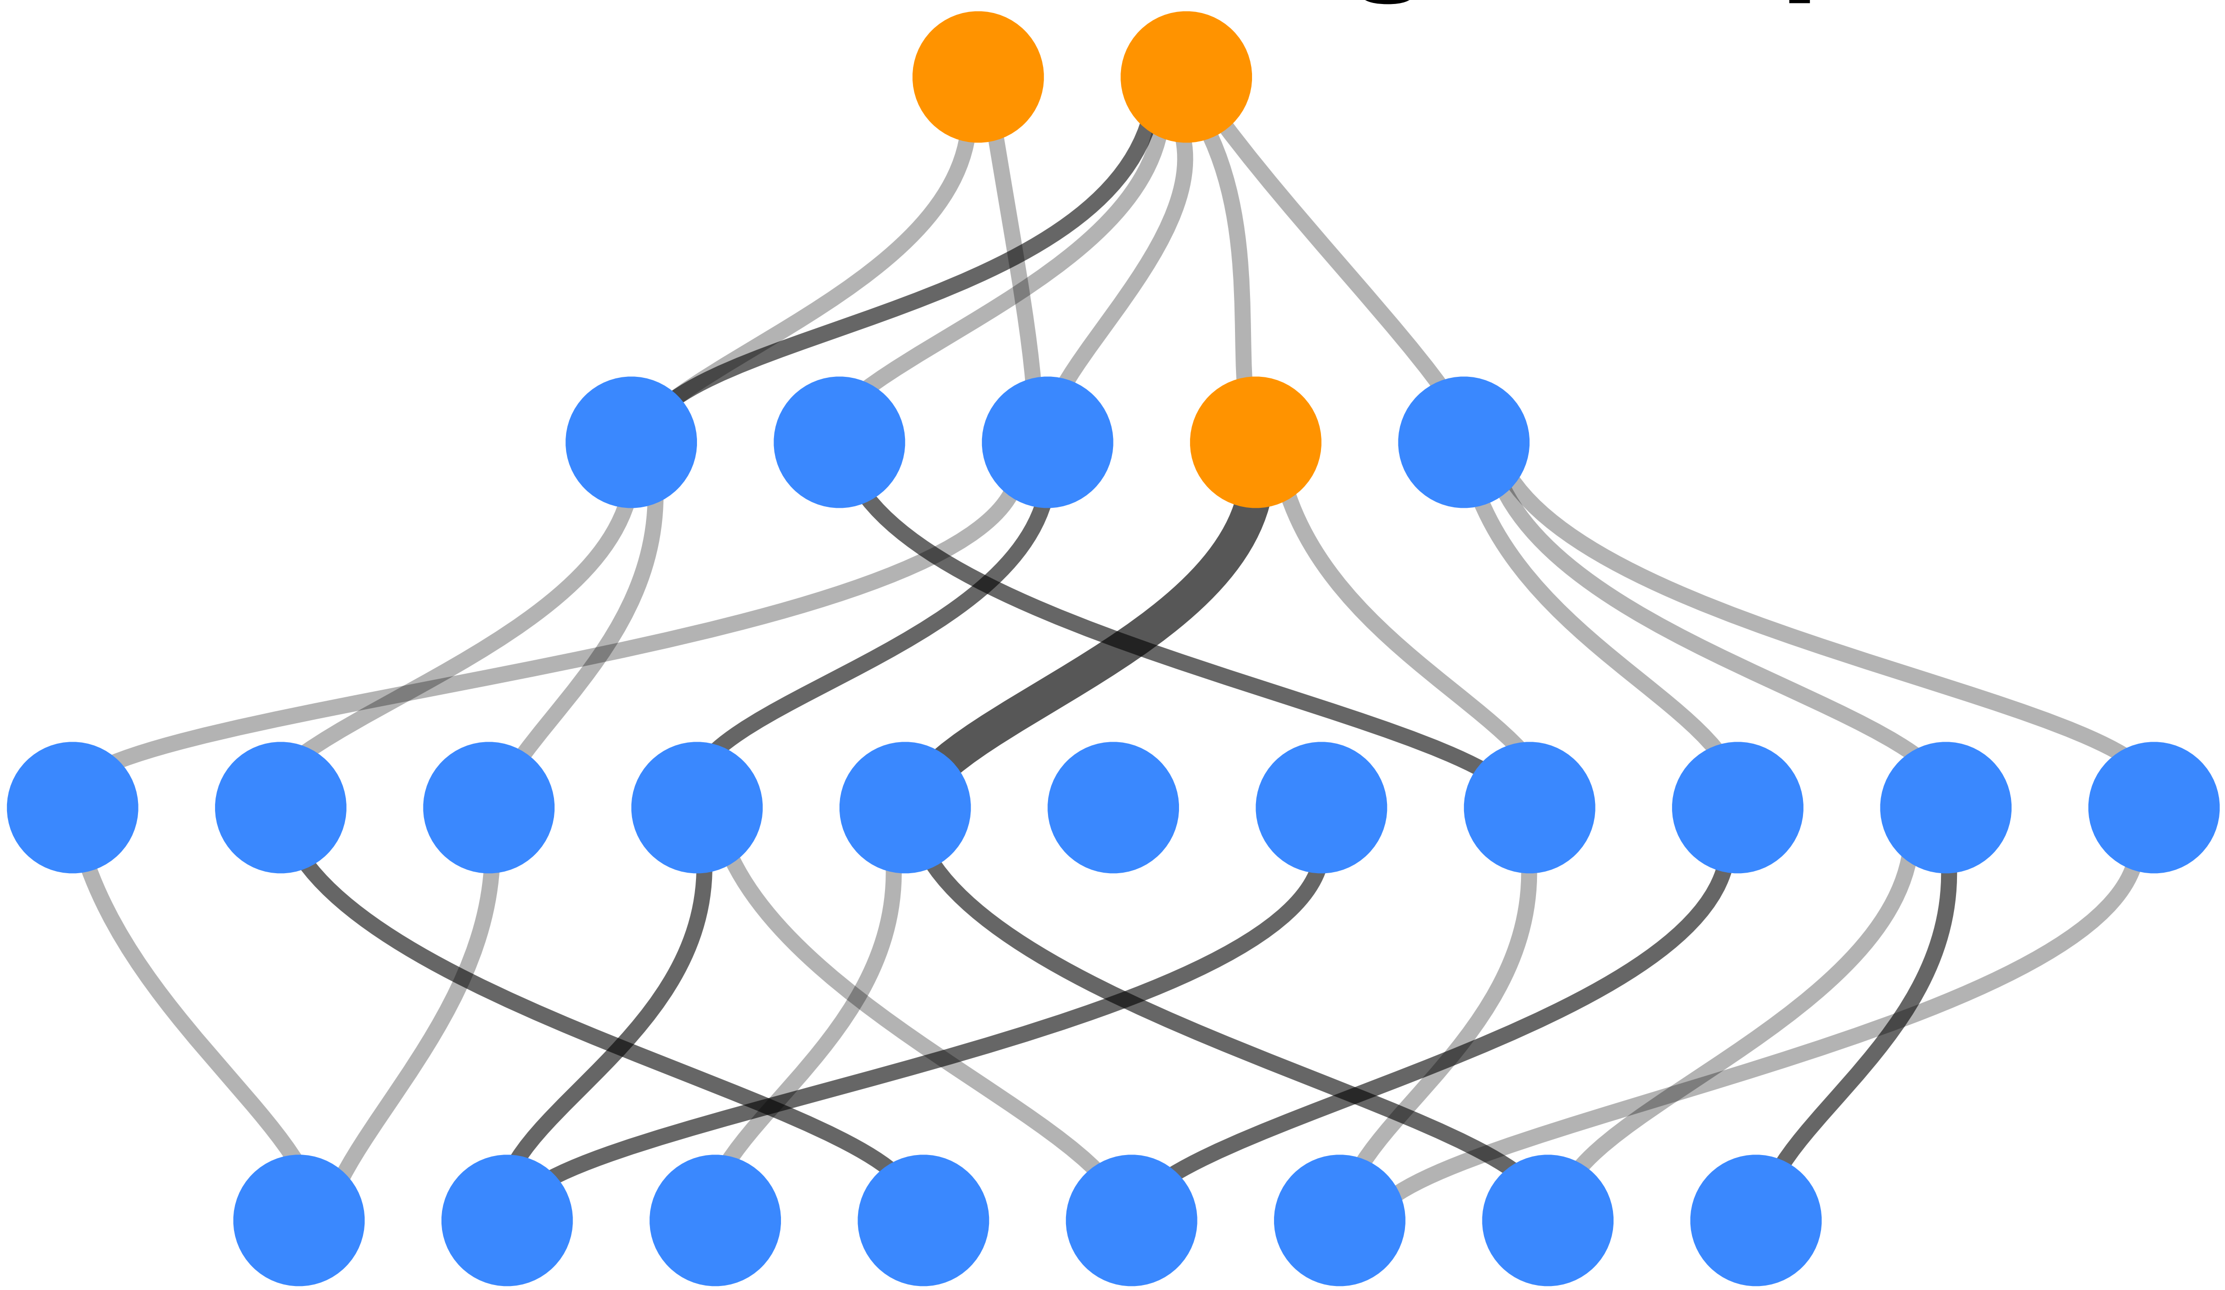

Machine learning is **not** a magic bullet!

Machine learning is **not** a magic bullet!

Training a neural network is a highly nonlinear and stochastic process (not well-understood theoretically)

Results depend on the quantity and quality of training data

It is no substitute for basic understanding

Interpretation requires care

Machine learning is **not** a magic bullet!

Training a neural network is a highly nonlinear and stochastic process (not well-understood theoretically)

Results depend on the quantity and quality of training data

It is no substitute for basic understanding

Interpretation requires care

…but it **can** be useful and is fun!
(much to explore)

Classical Machine Learning for Quantum Technology – a brief survey of an evolving field

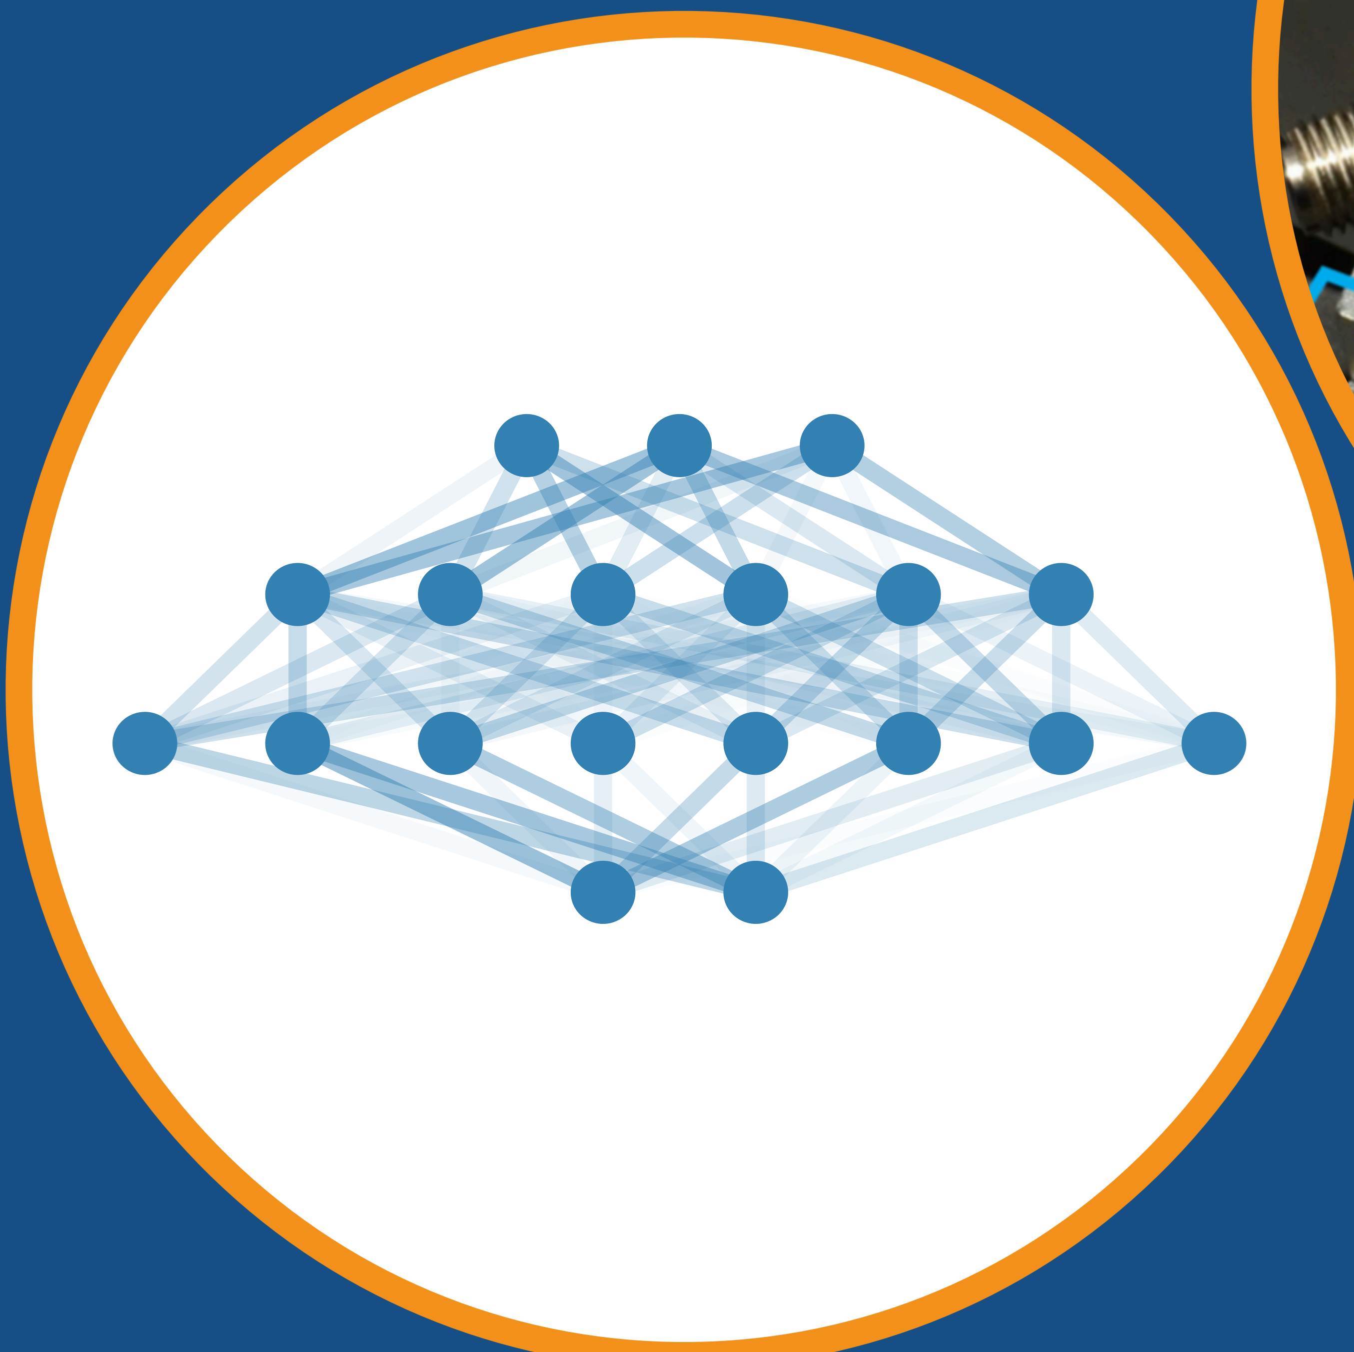

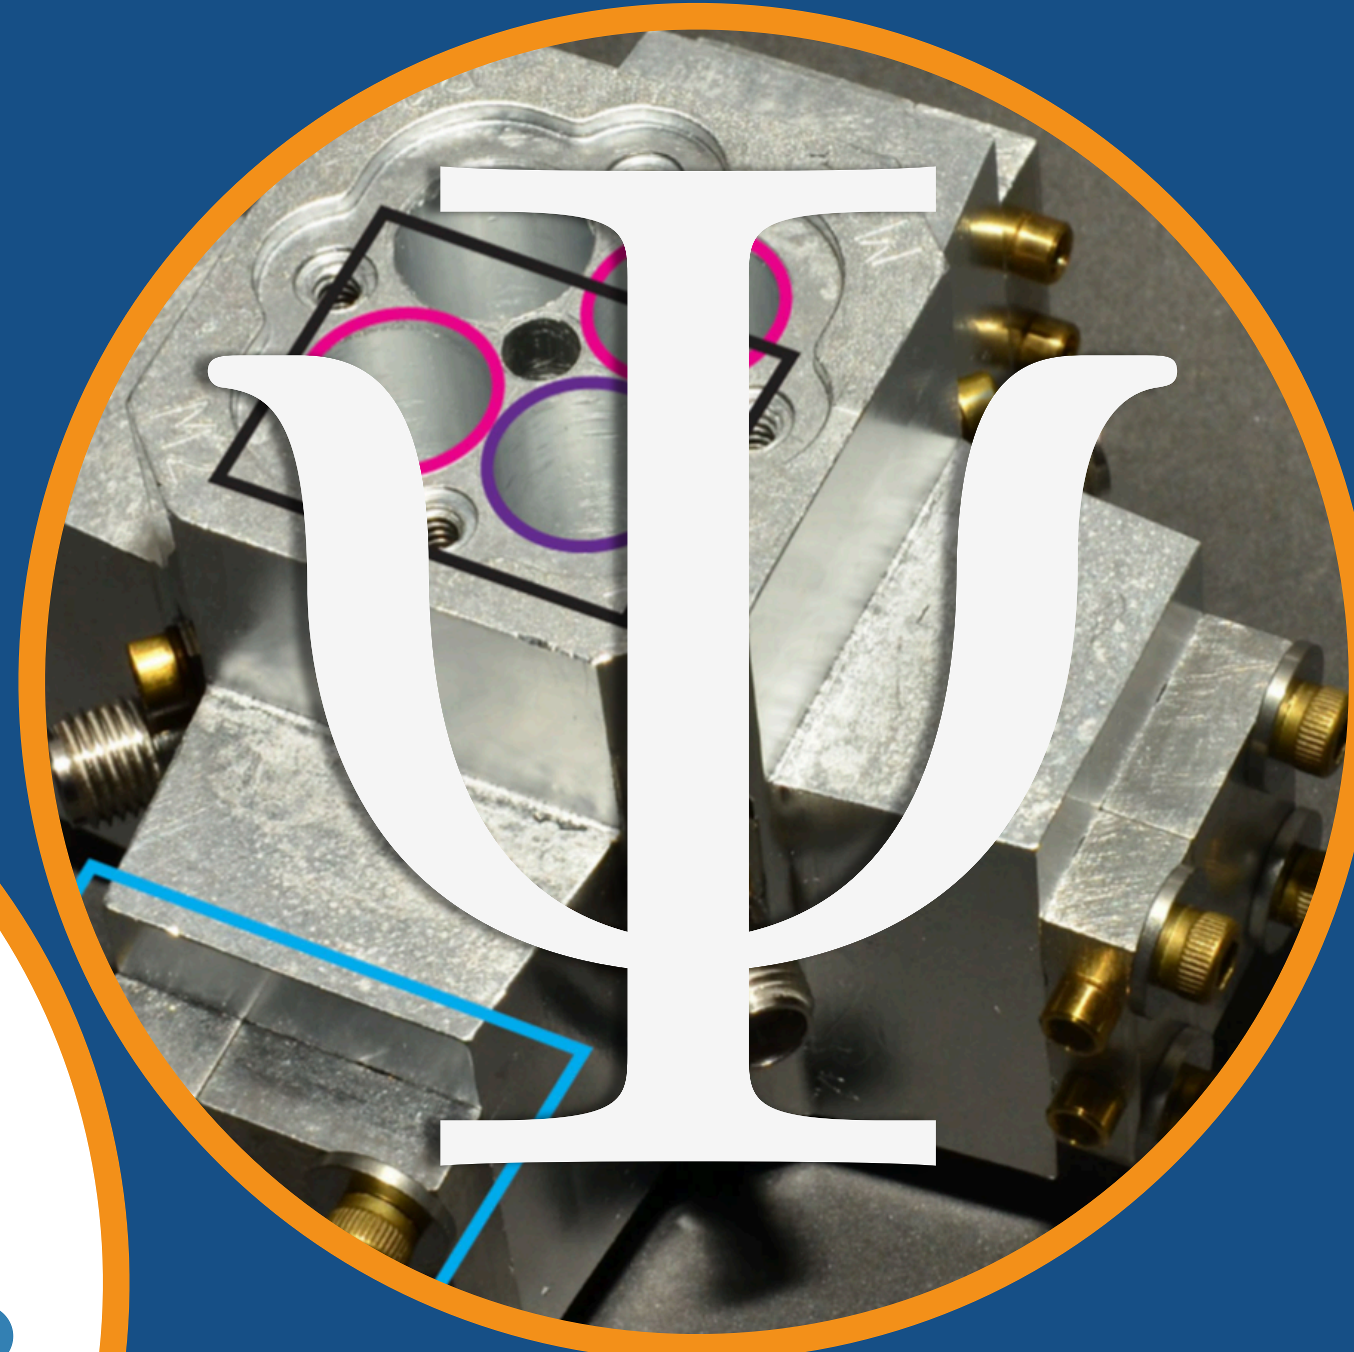

image: Chou et al. Nature 2018

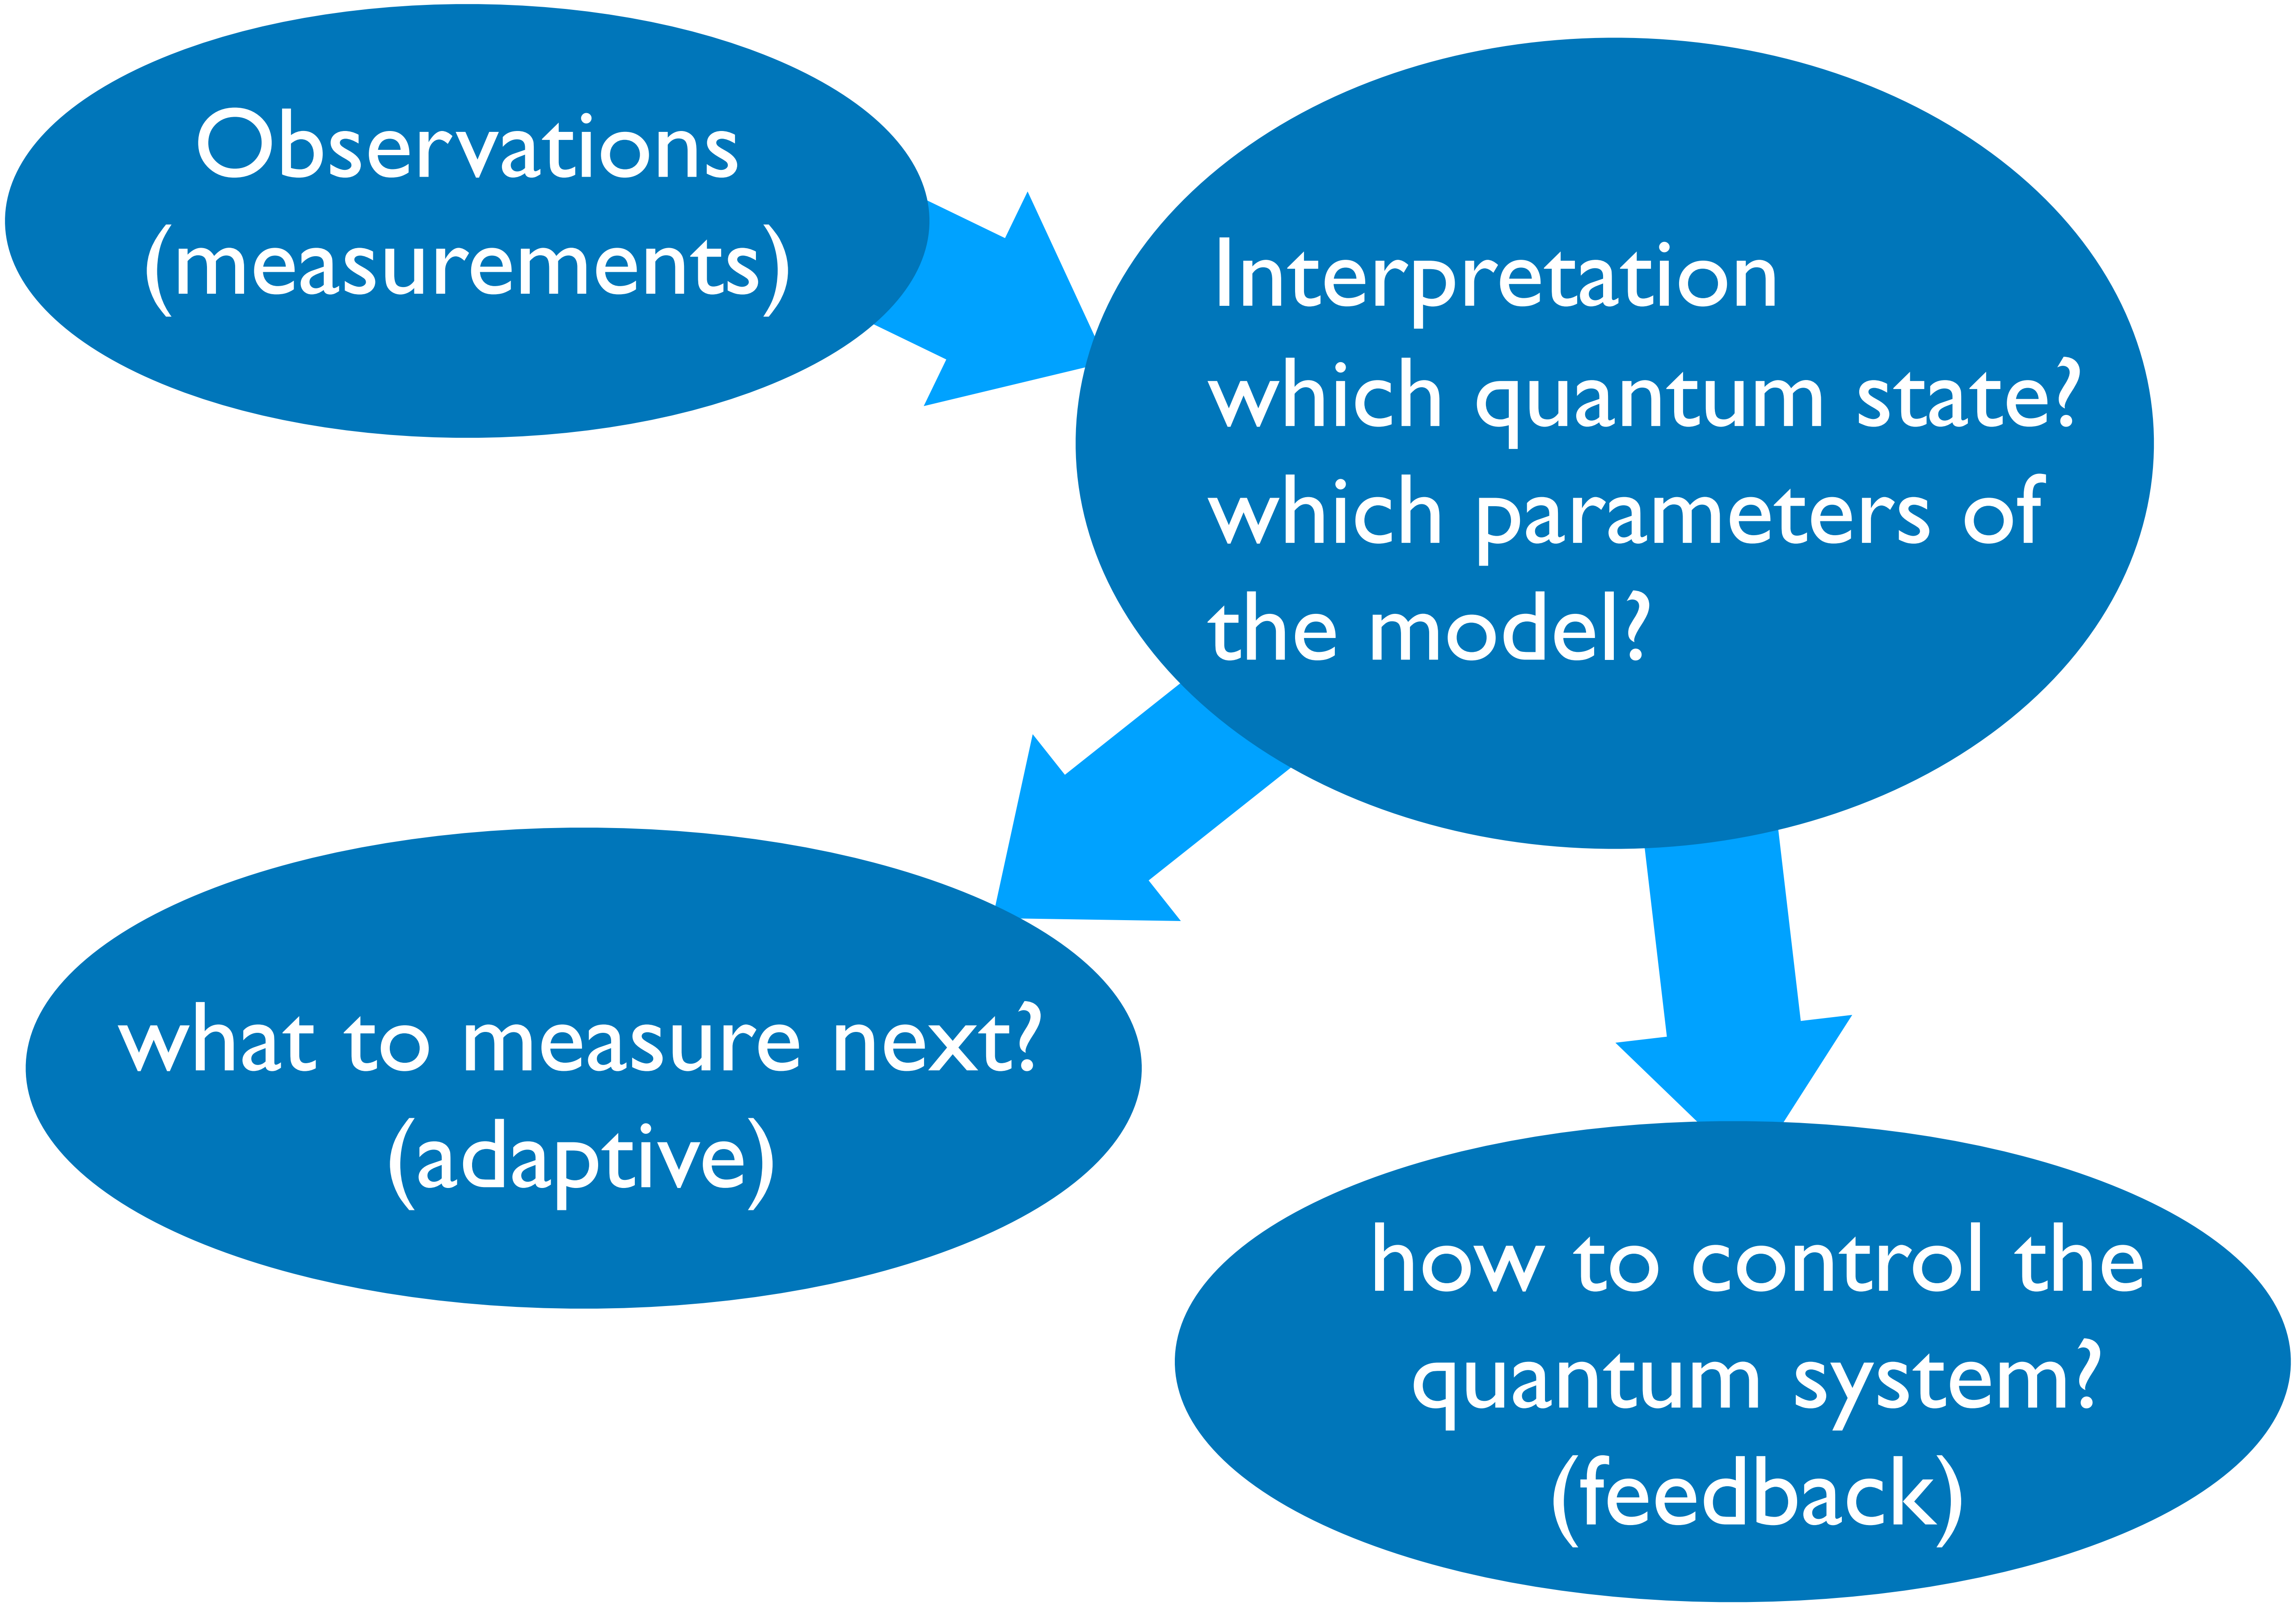
Observations
(measurements)
Interpretation
which quantum state?
which parameters of
the model?
what to measure next?
(adaptive)
how to control the
quantum system?
(feedback)

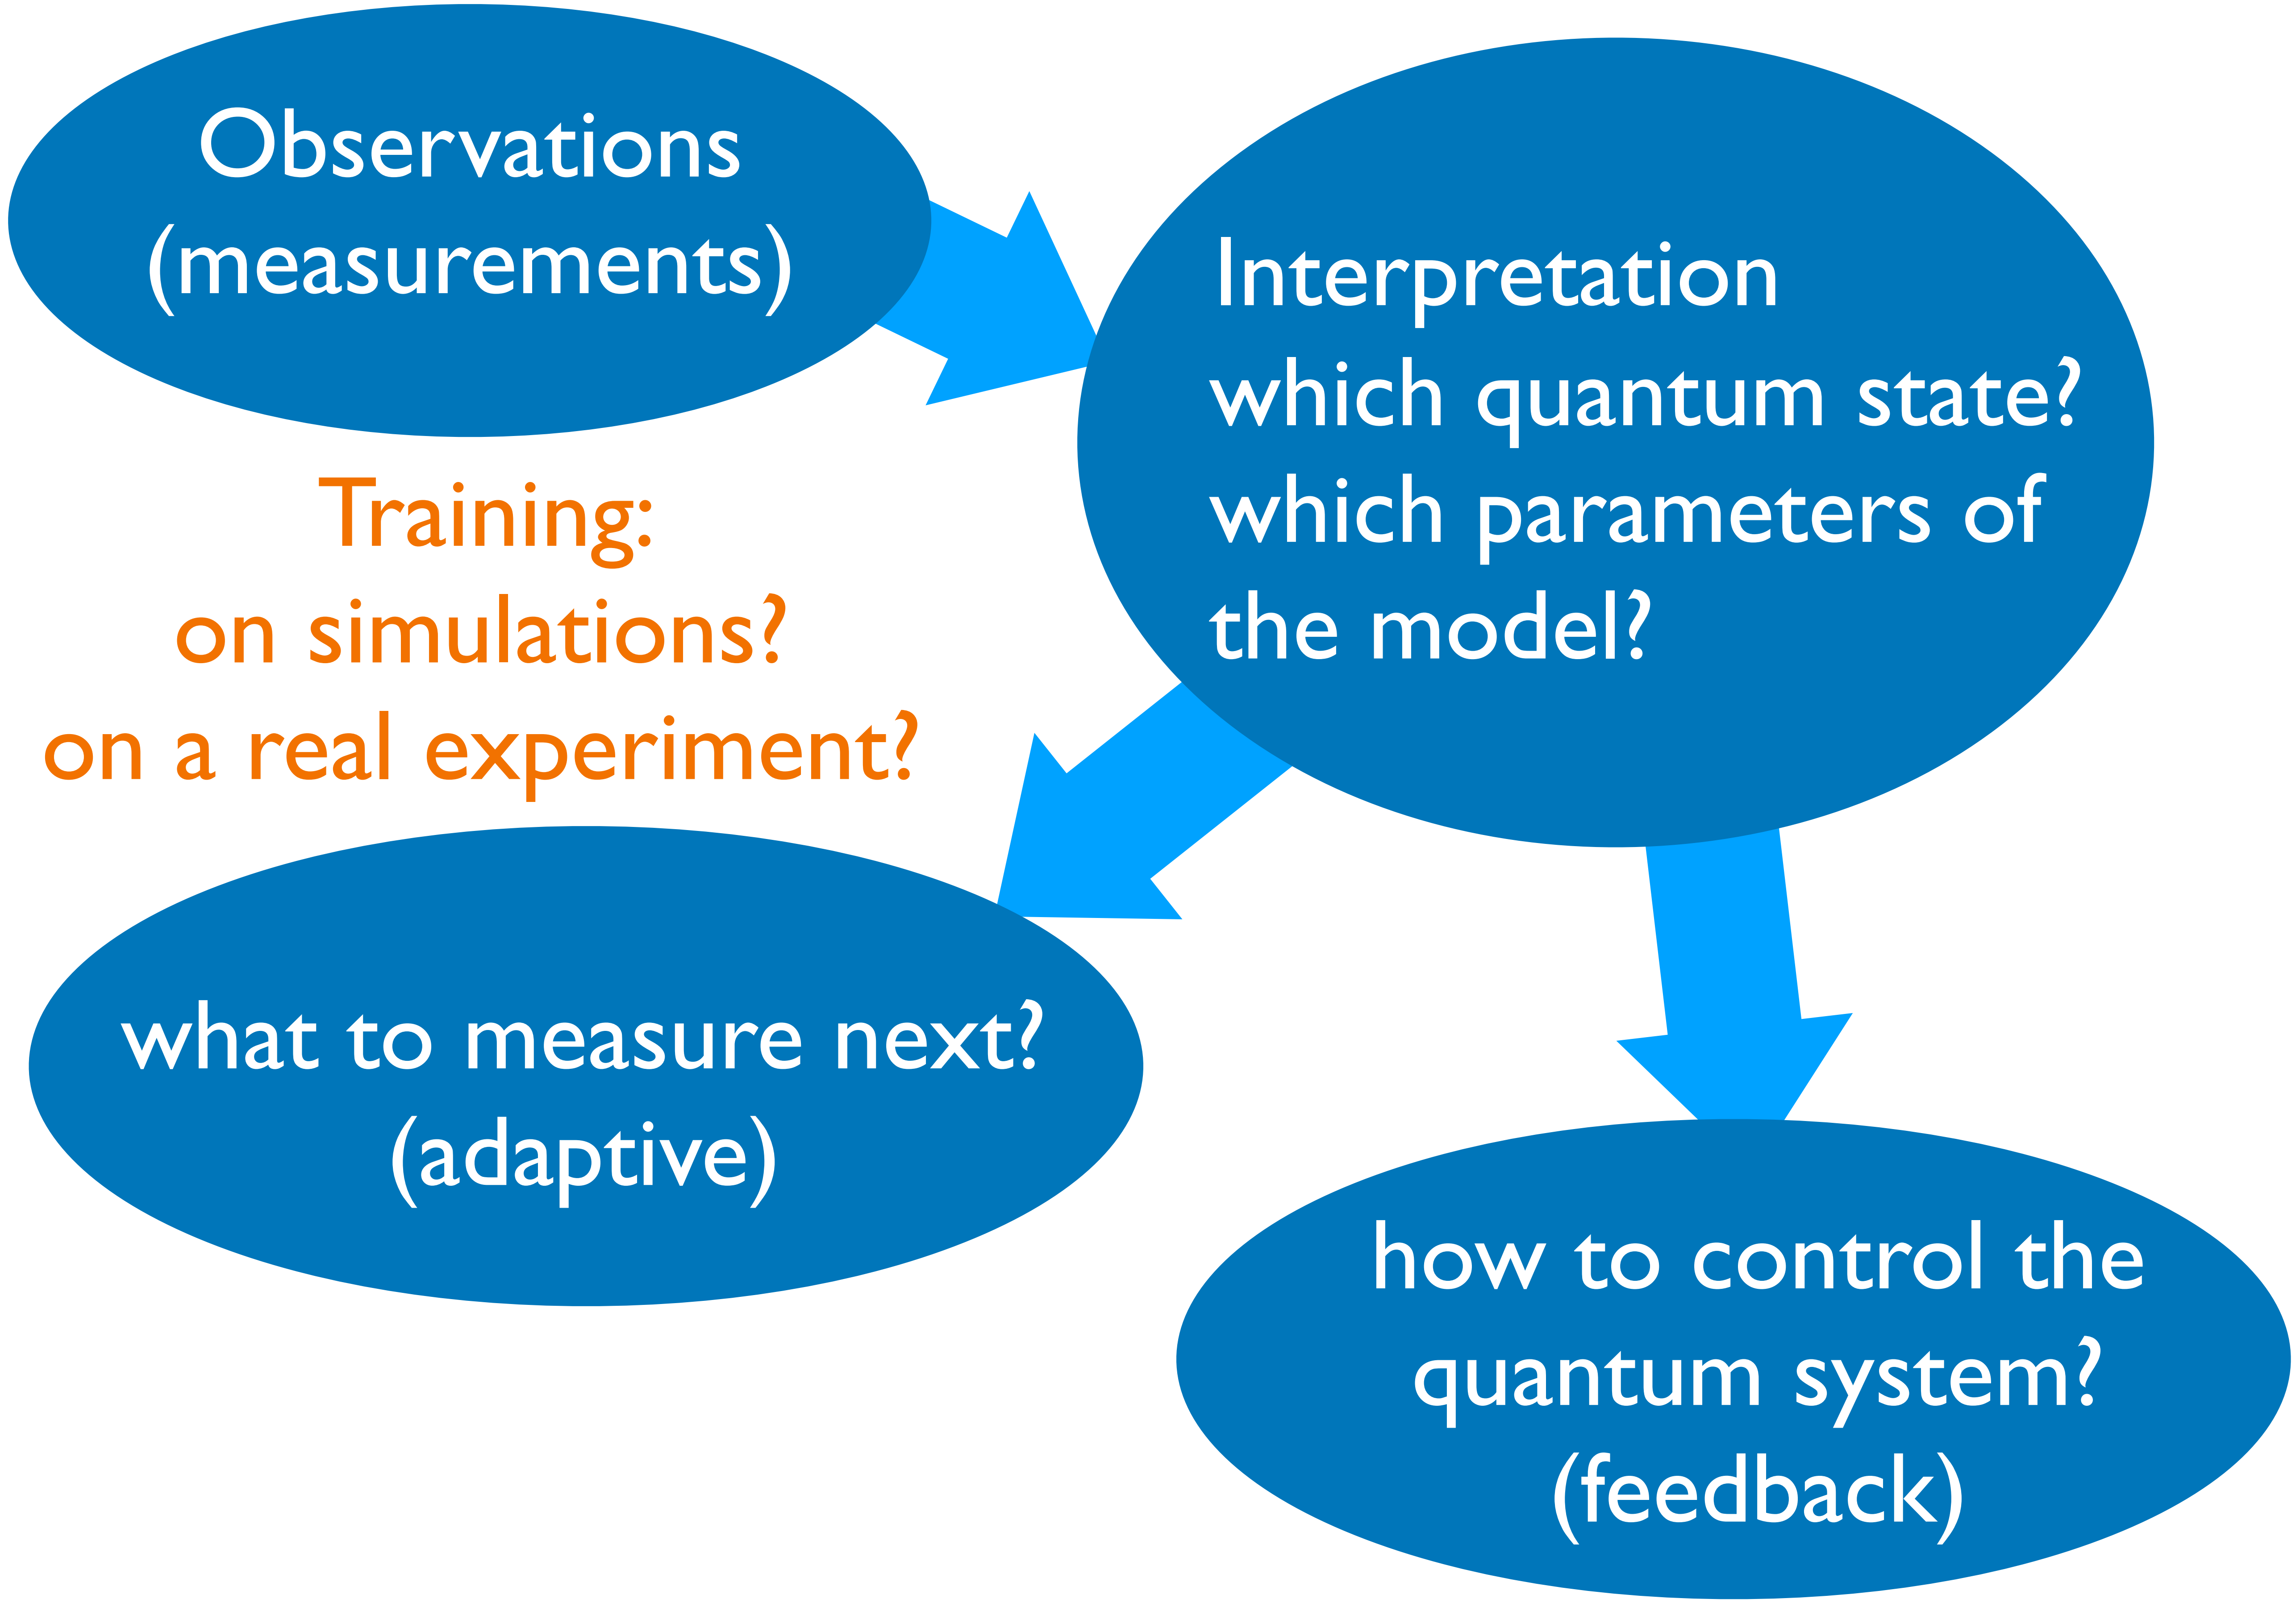
Observations
(measurements)
Interpretation
which quantum state?
which parameters of
the model?
Training:
on simulations?
on a real experiment?
what to measure next?
(adaptive)
how to control the
quantum system?
(feedback)

Readout: time traces

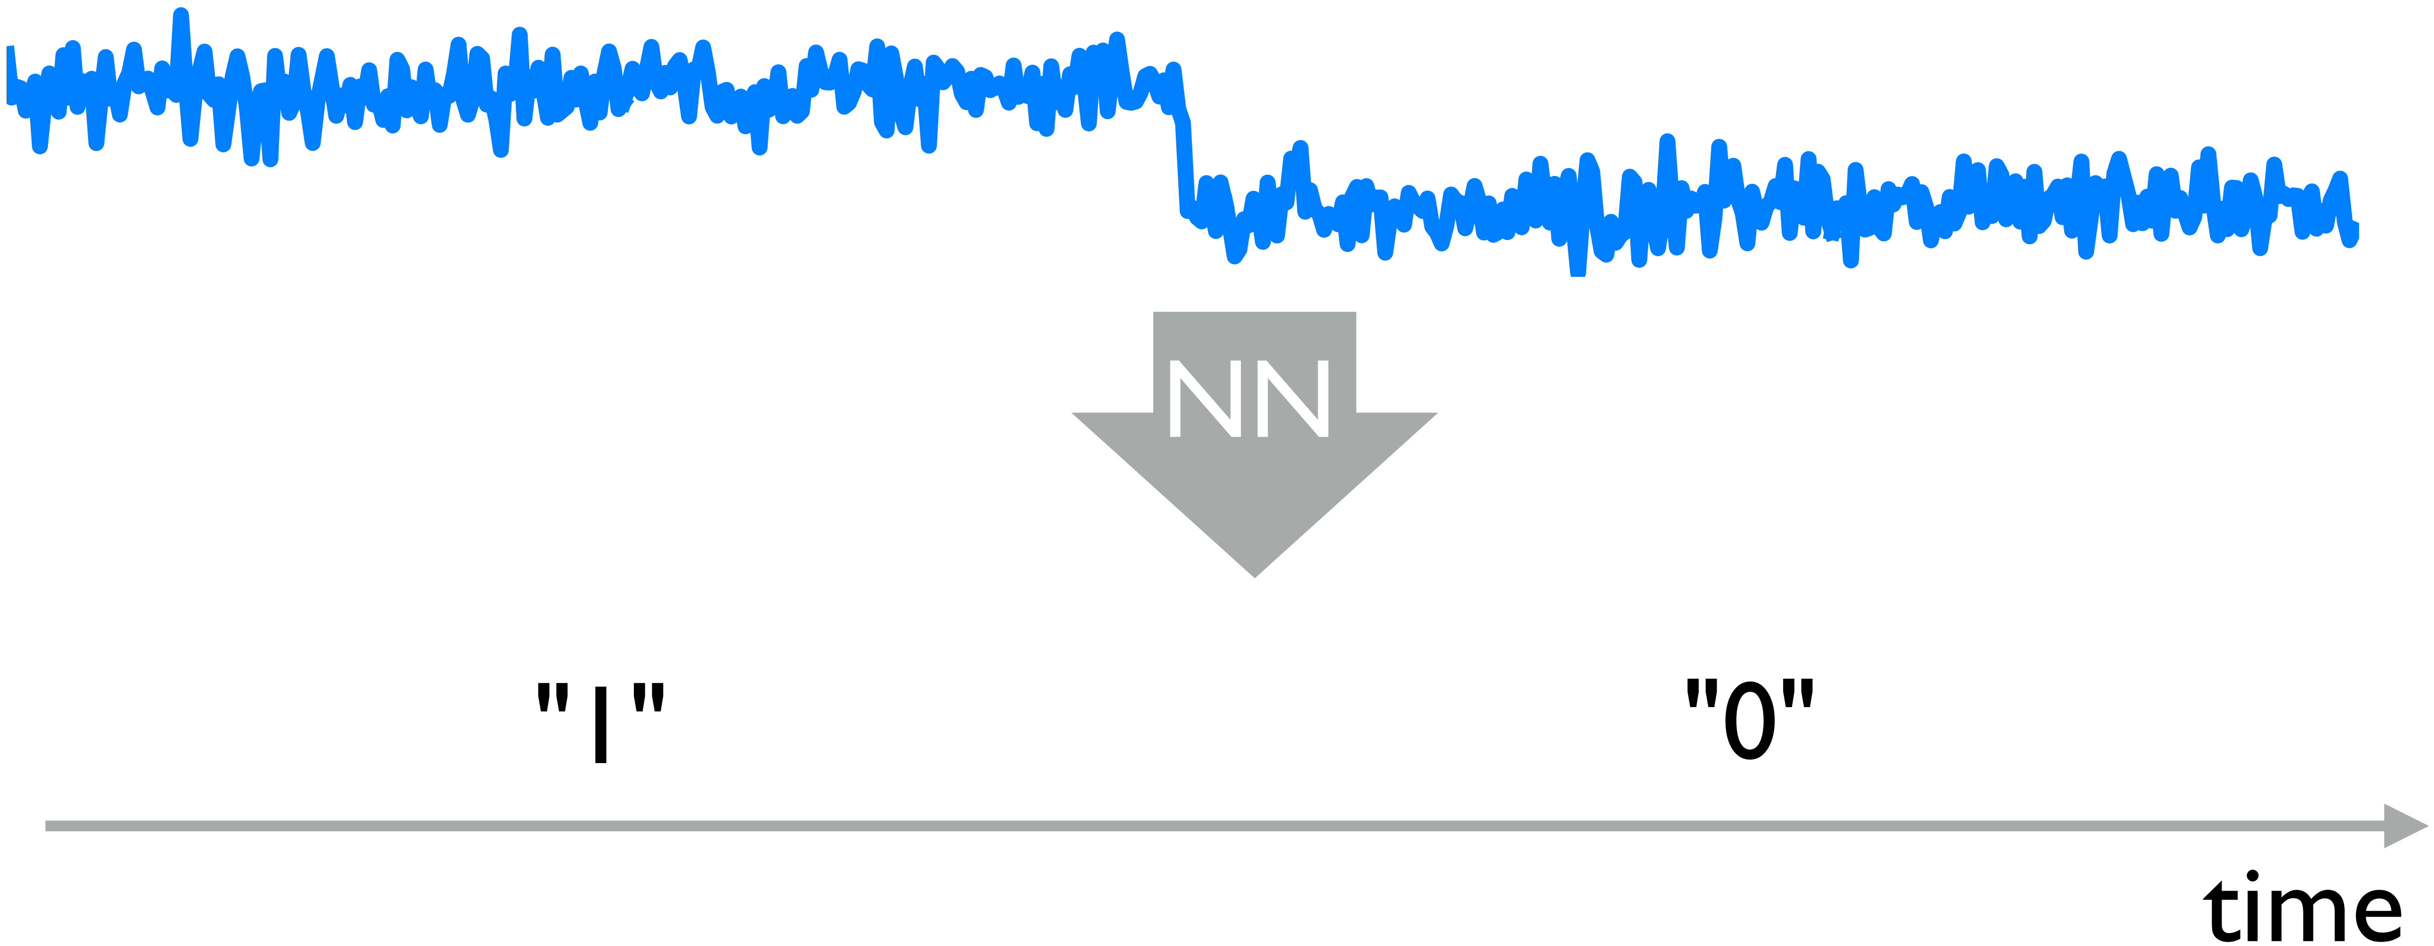

(use recurrent neural networks, i.e. nets with memory)

advantage: NN trained on real data learns all possible distortions and noise sources

Readout: time traces

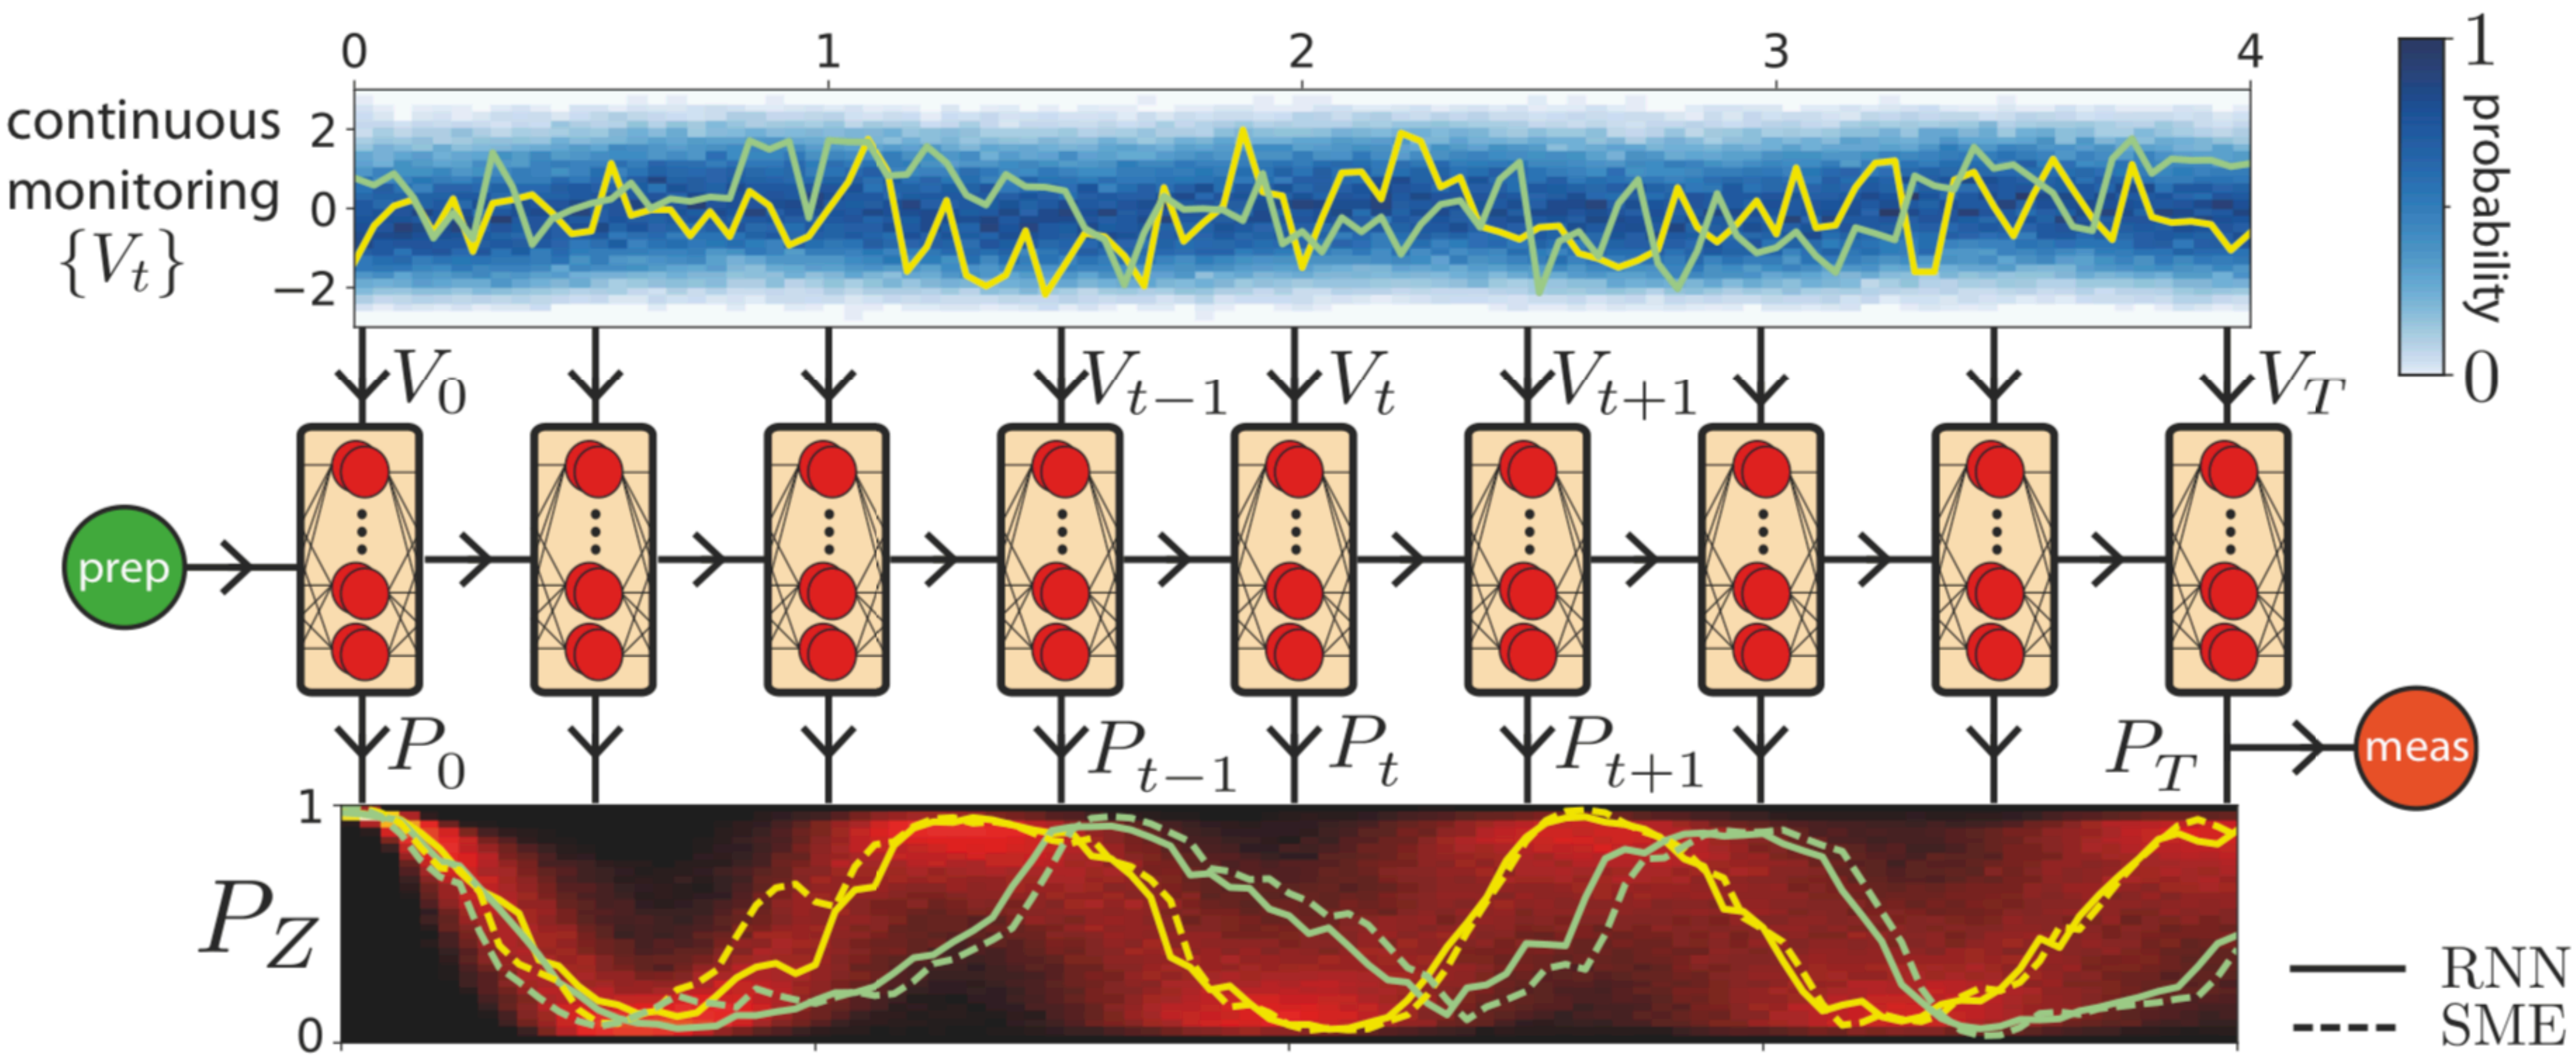

E. Fleurin et al 2018 (s.c. qubit readout experiment)

Interpreting error syndromes

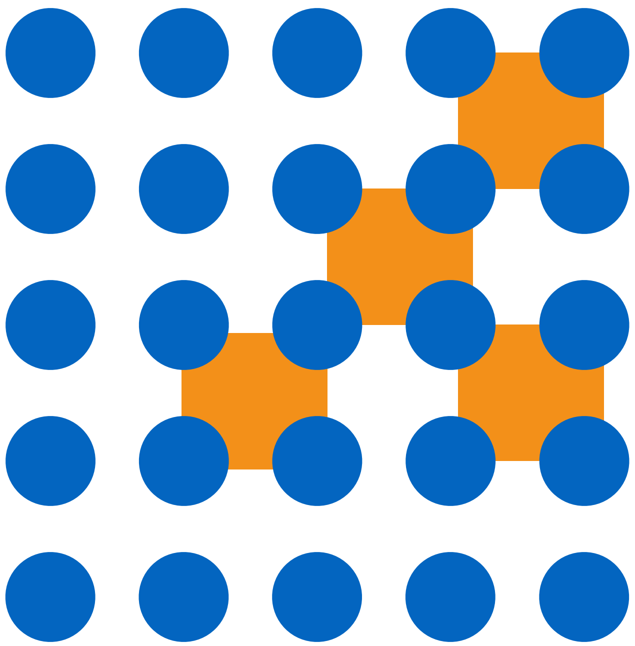

(surface code)

Interpreting error syndromes

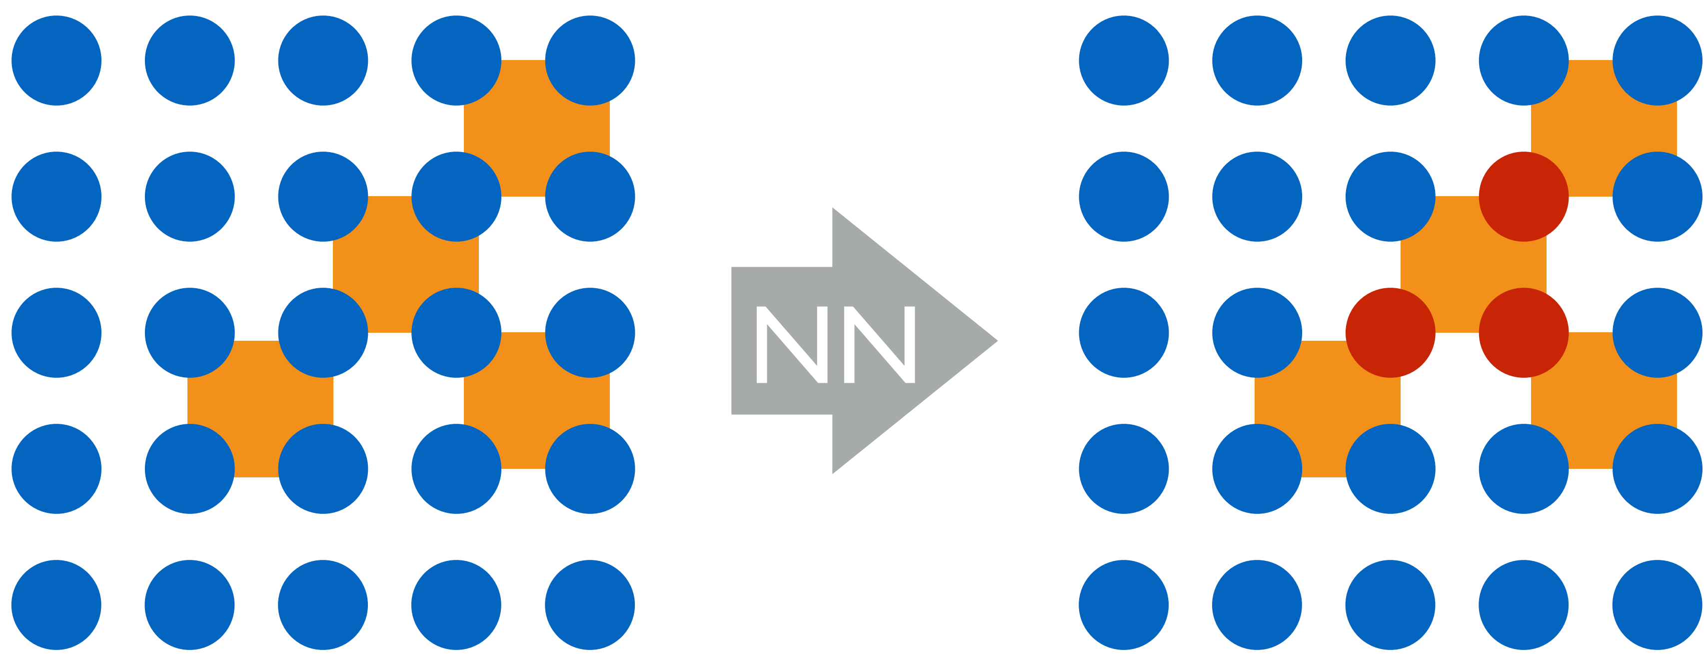

(surface code)

Melko et al 2017, Jiang et al, Bayreuther and Beenakker, …

Adaptive measurements

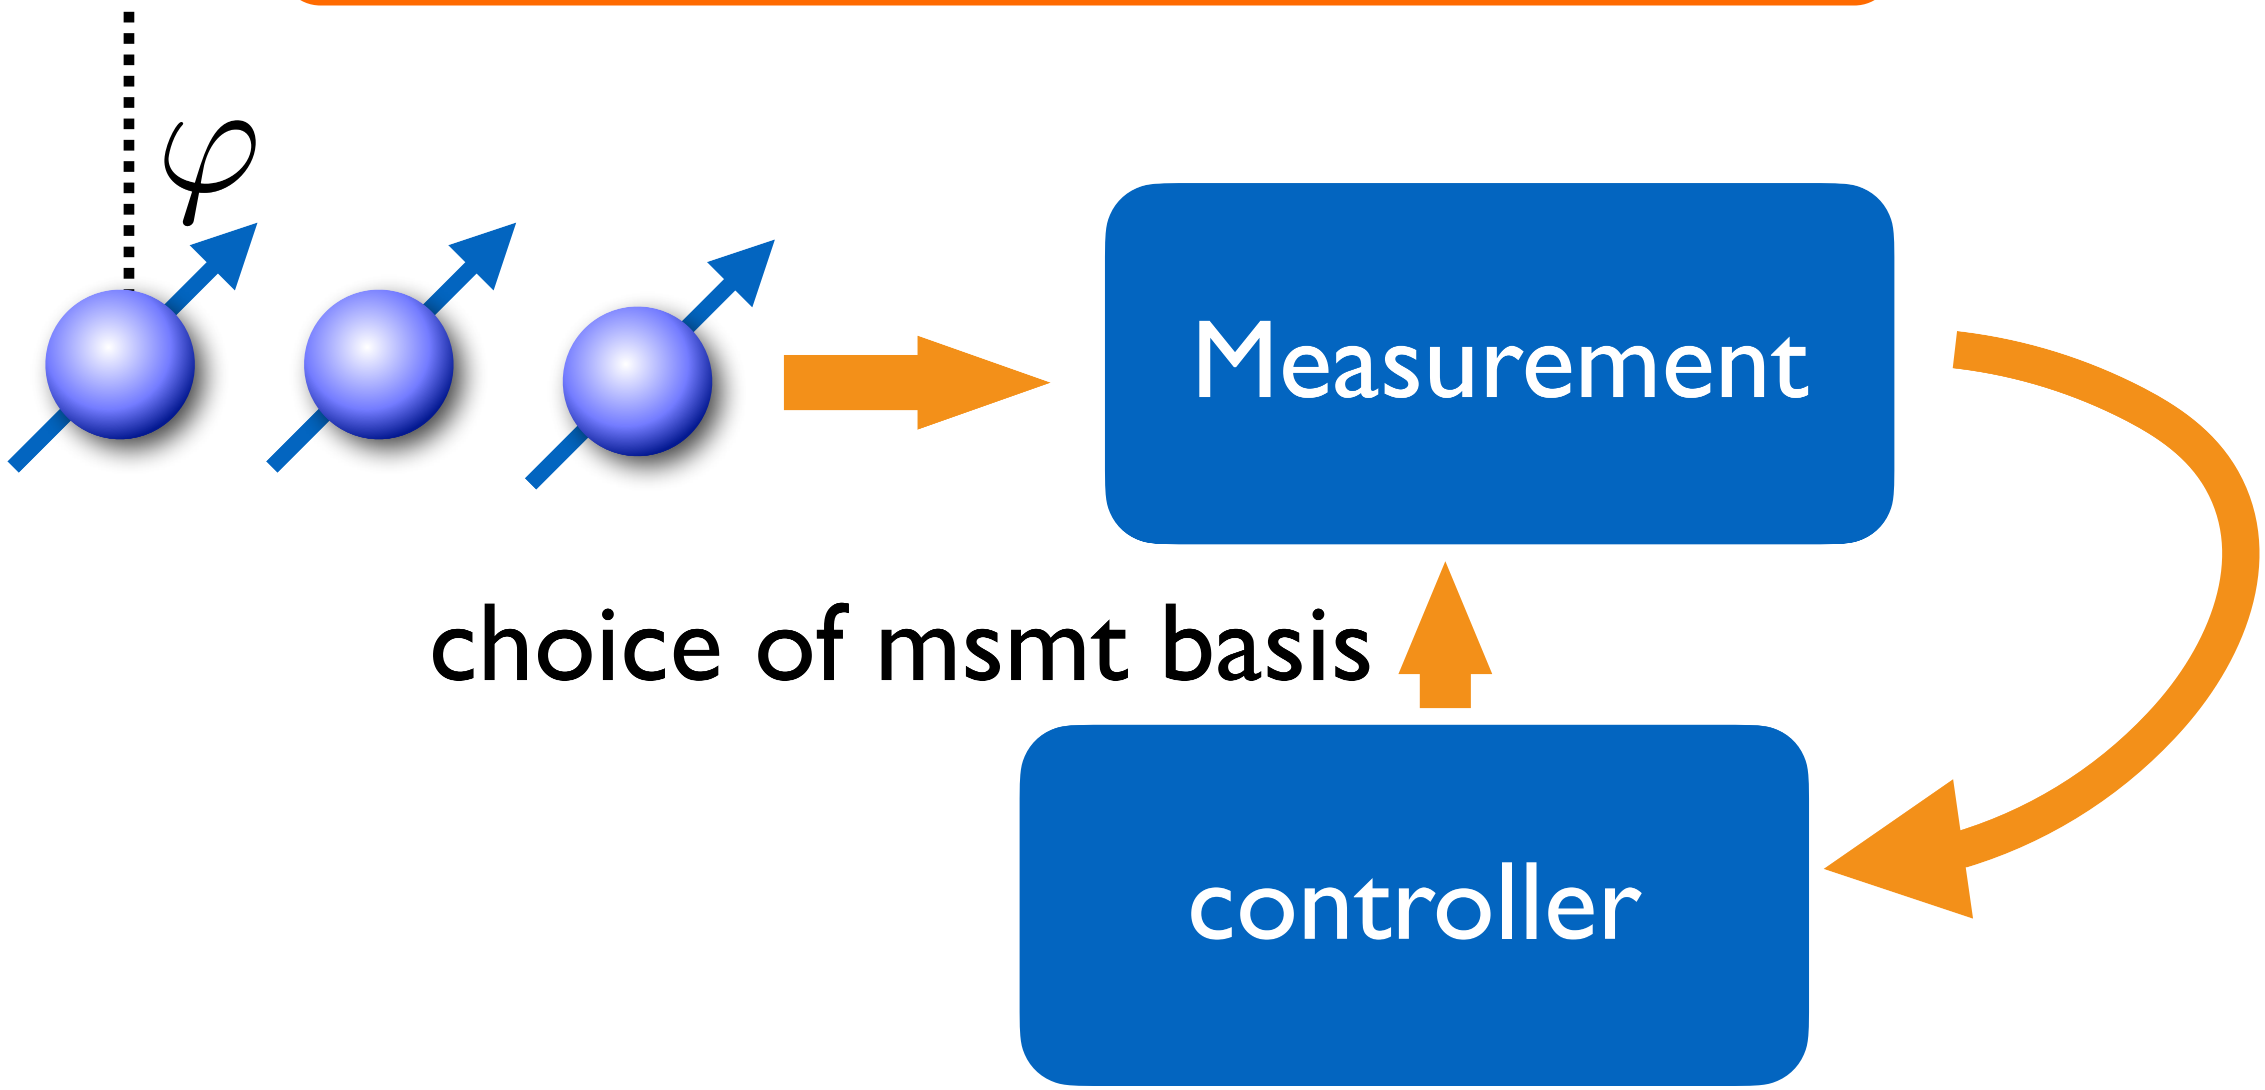

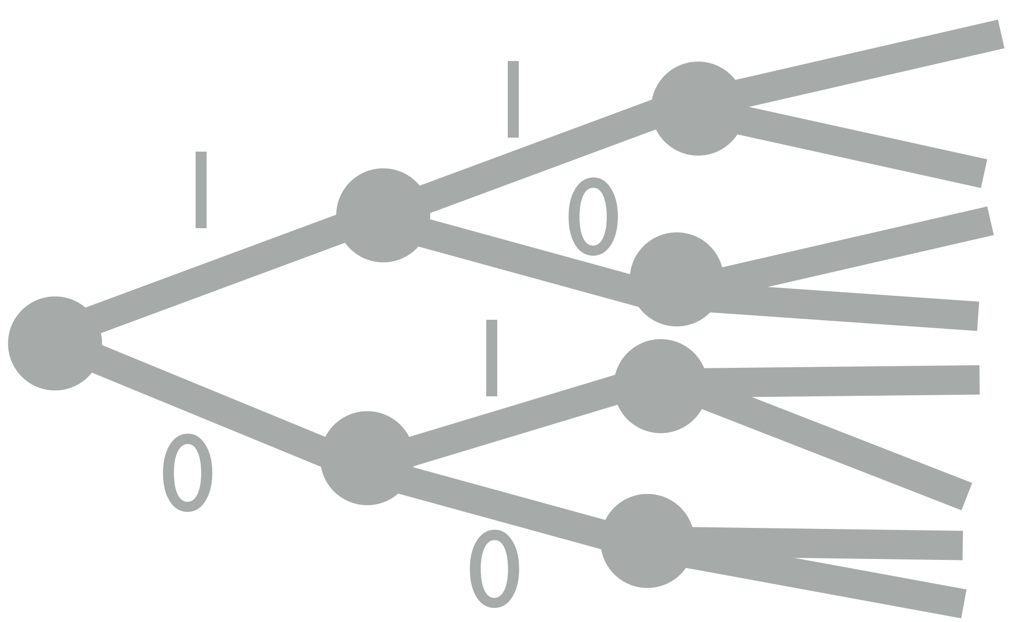

B. Sanders group 2010-now (particle swarm)

with NN: e.g. Quek et al 2021

Reconstruction of quantum states

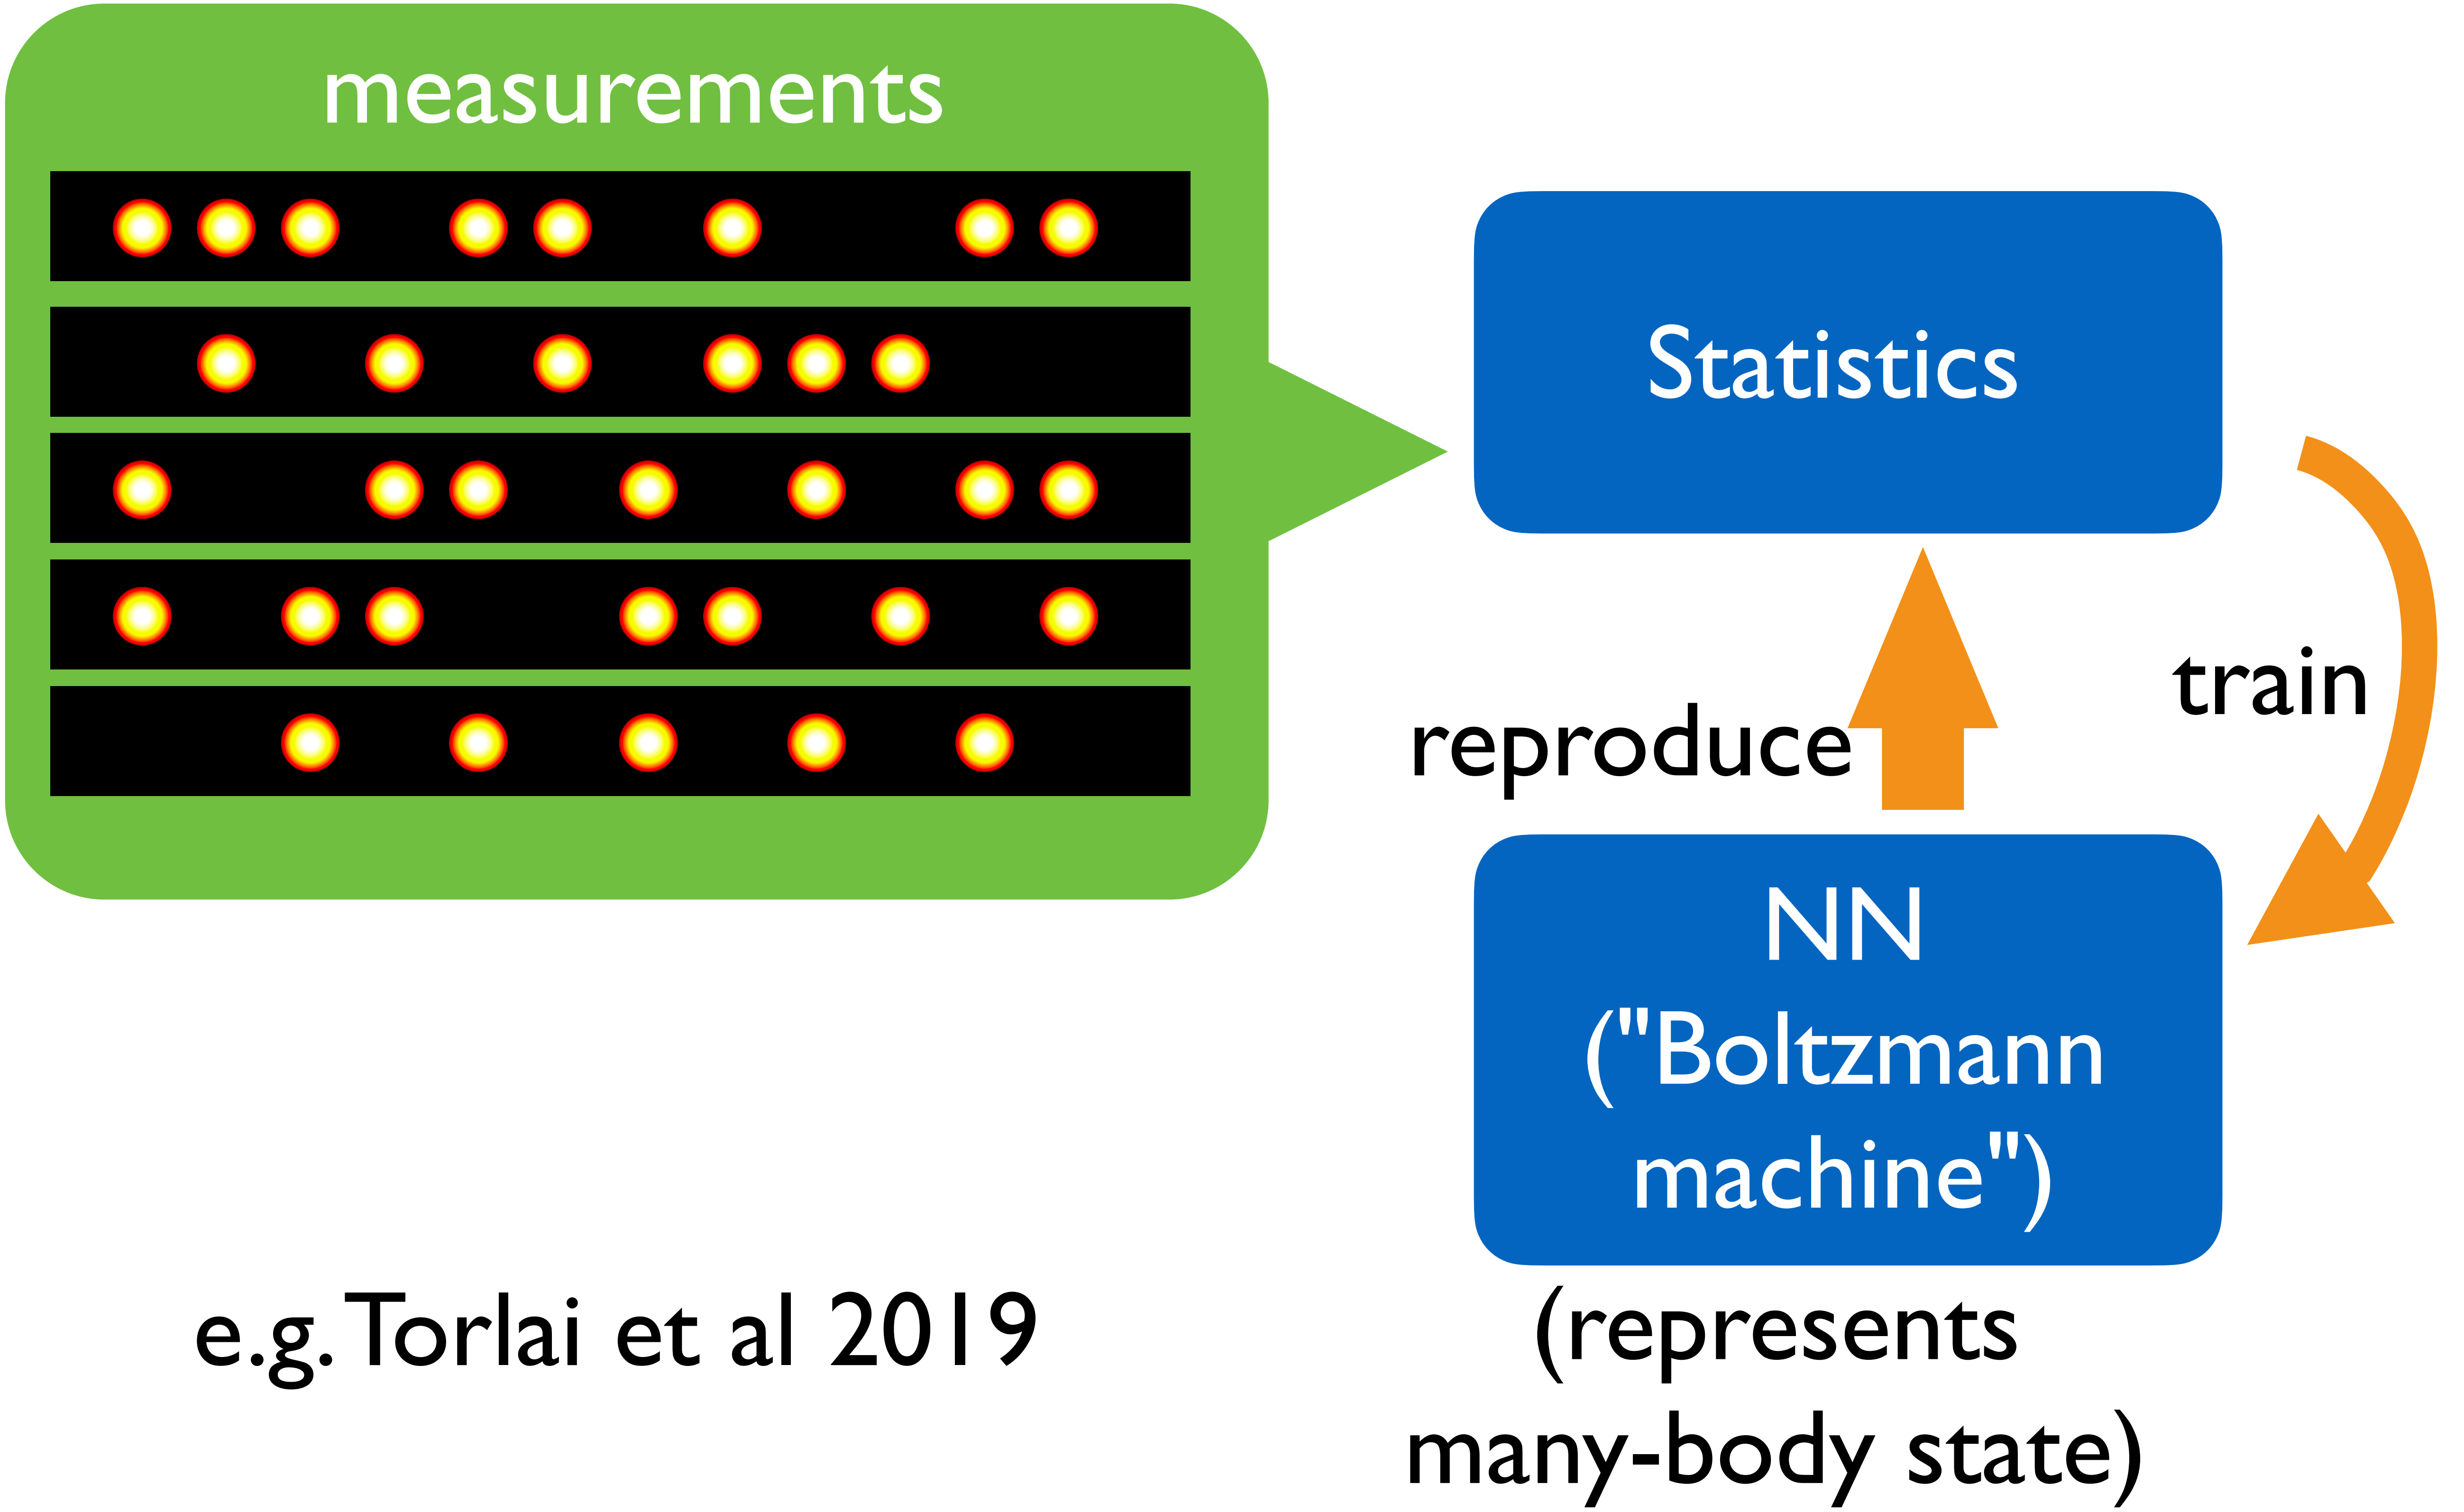

e.g. Torlai et al 2019

Calibration/tuning/learning Hamiltonians

$$\hat{H} = ?$$

Calibration/tuning/learning Hamiltonians

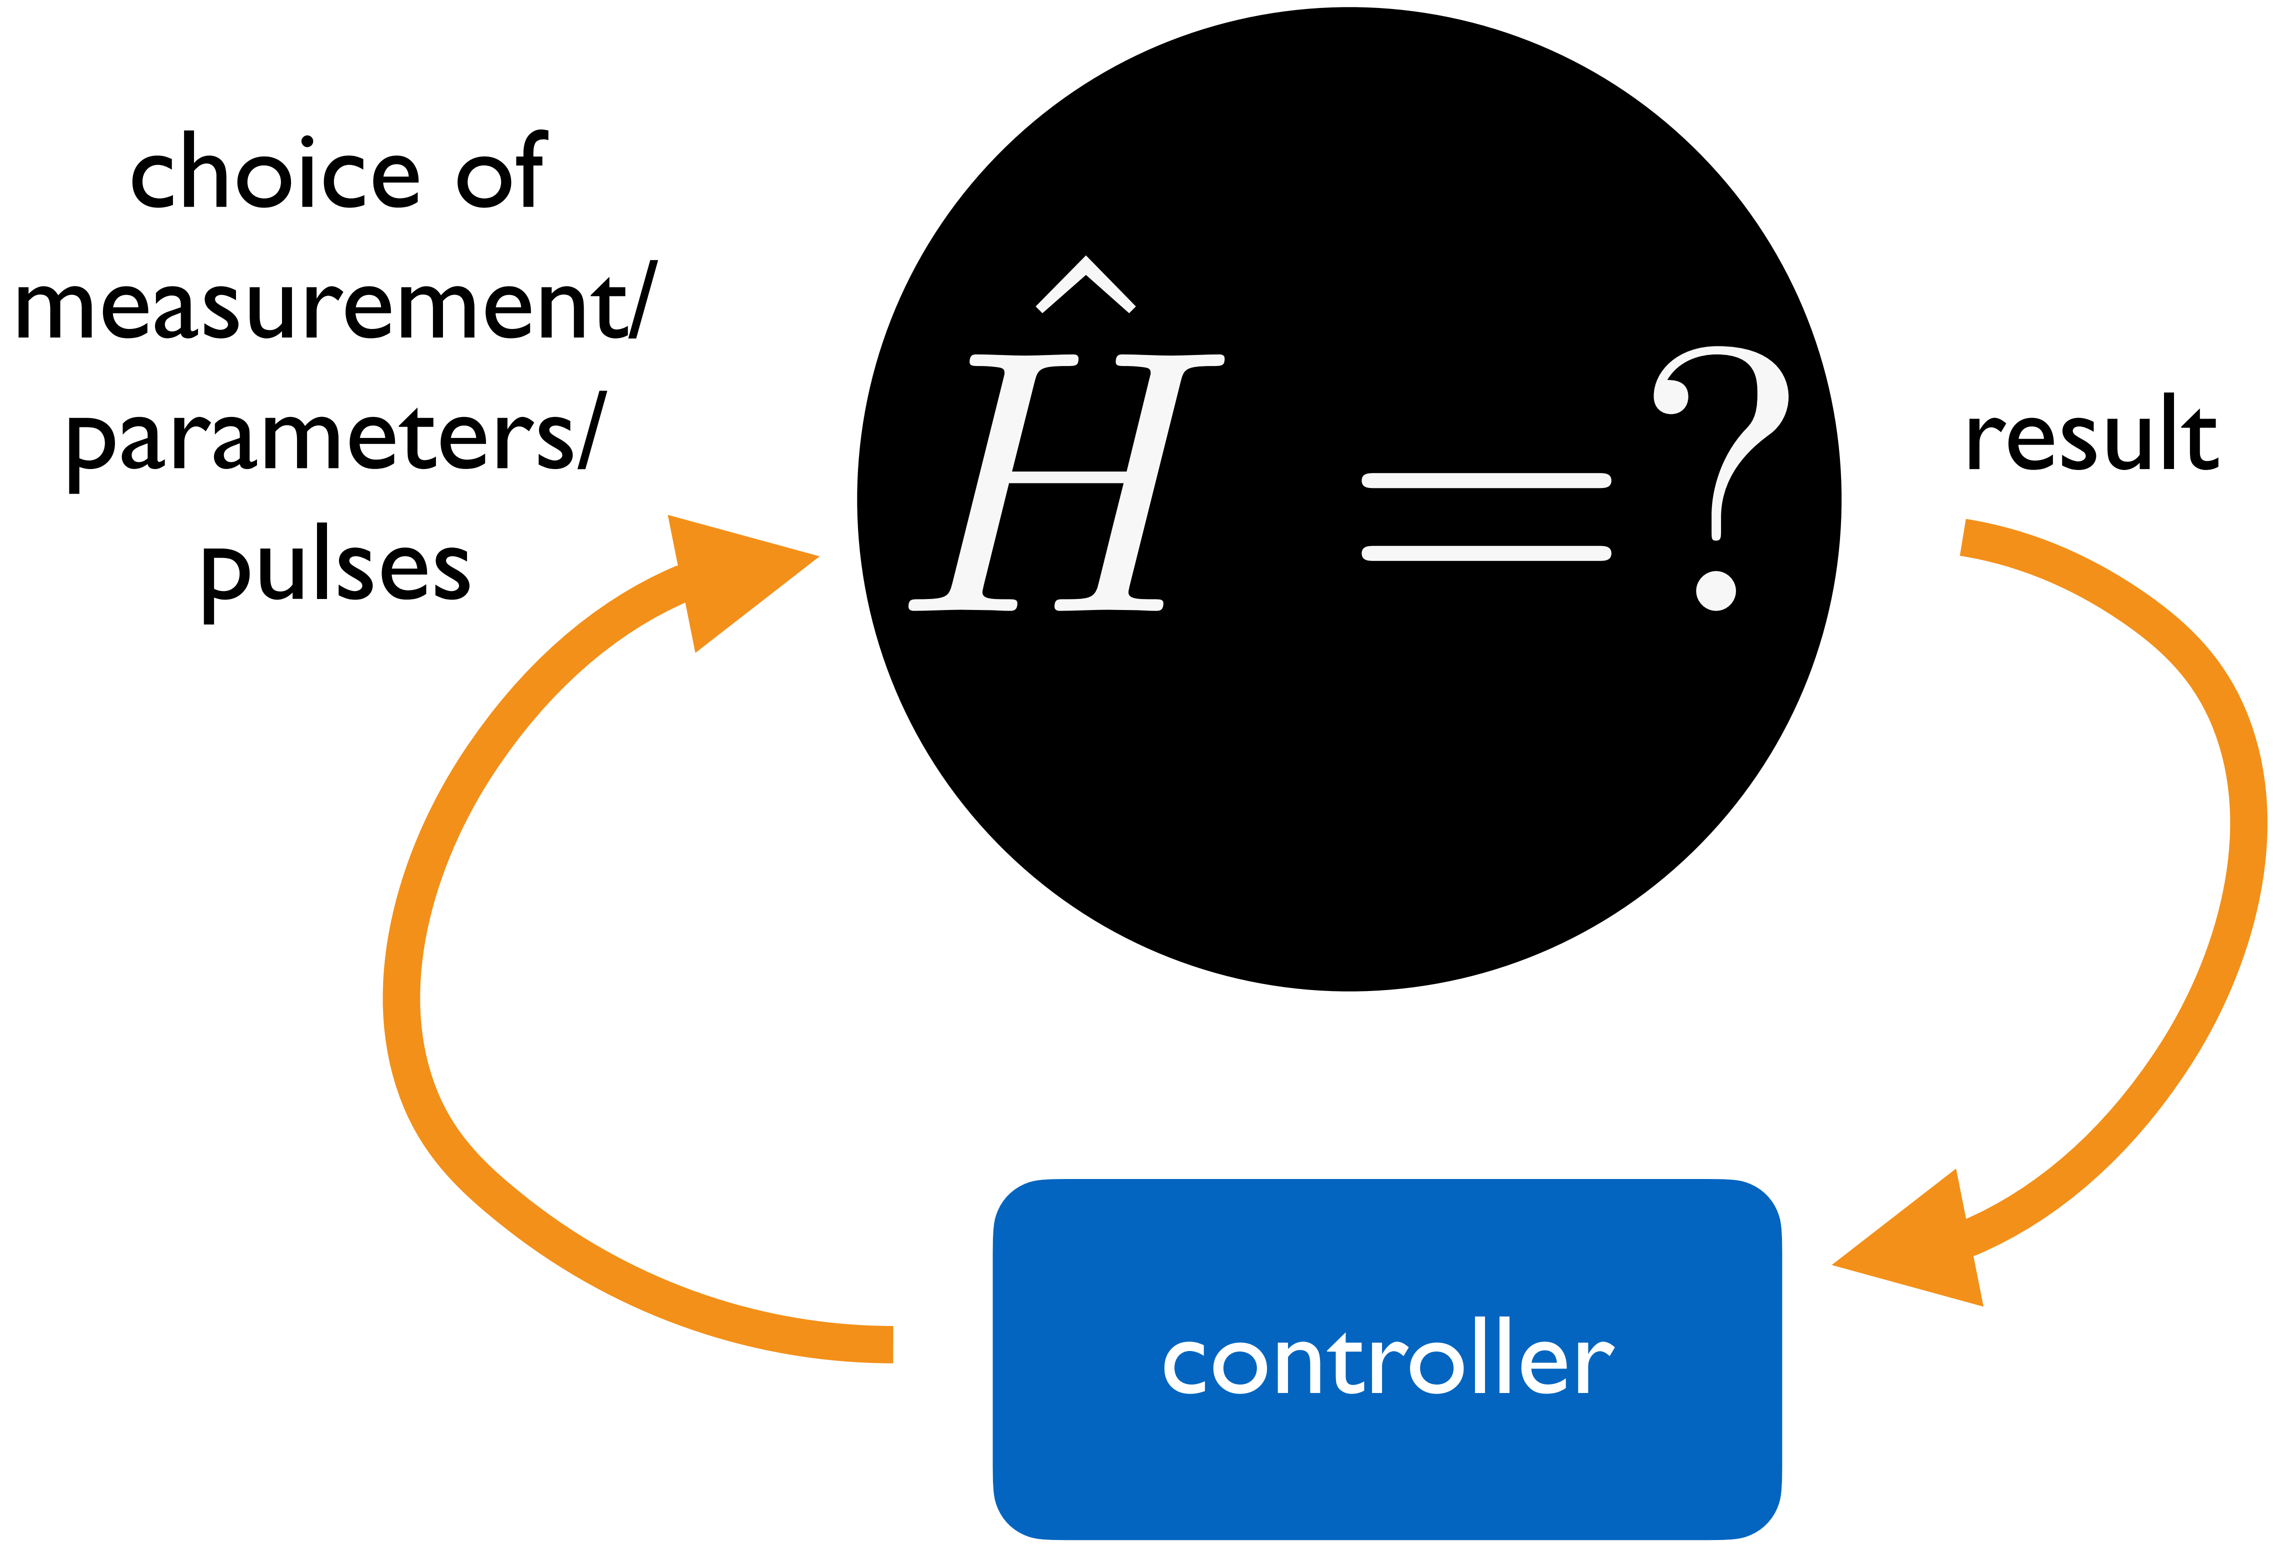

Calibration/tuning/learning Hamiltonians

measuring a quantum dot:
deciding where to measure next

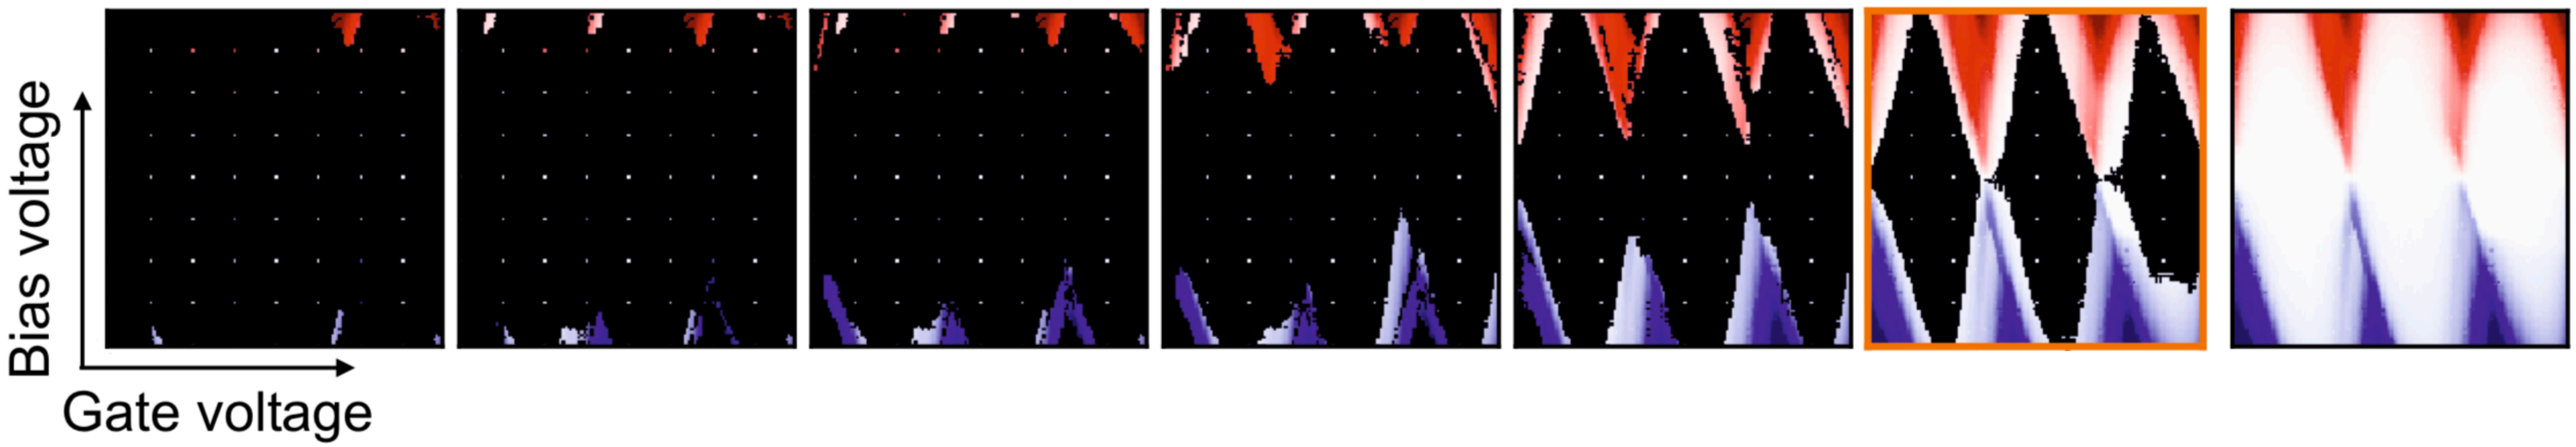

Natalia Ares group, 2018 & …

Calibration/tuning/learning Hamiltonians

Setup for diamond NV centre

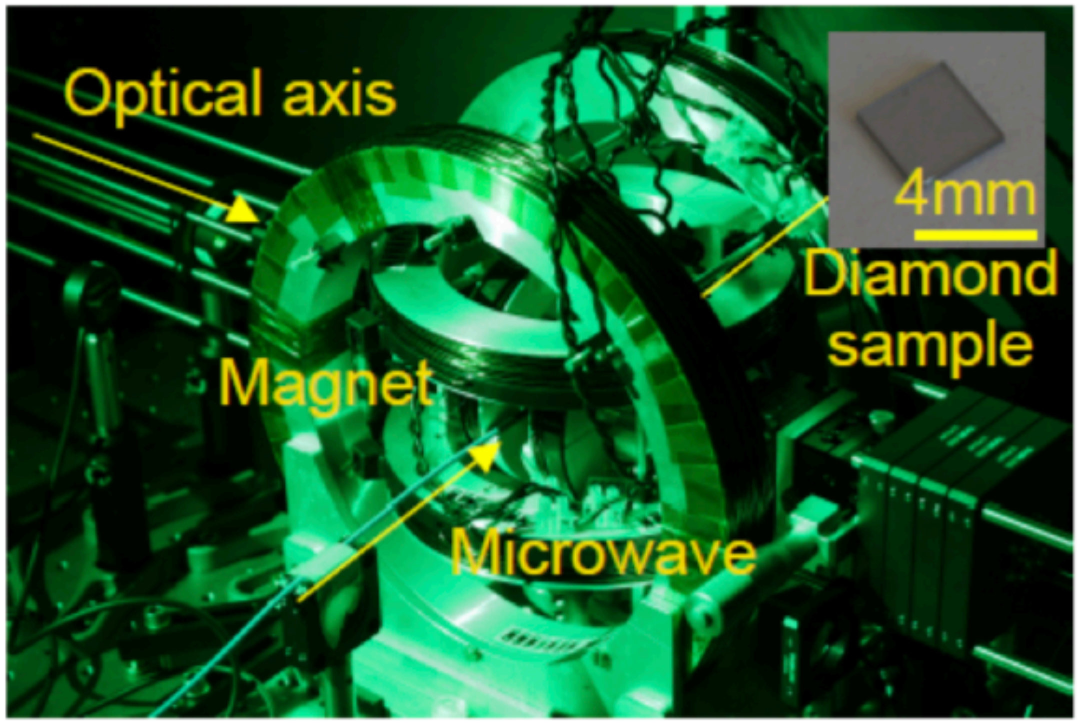

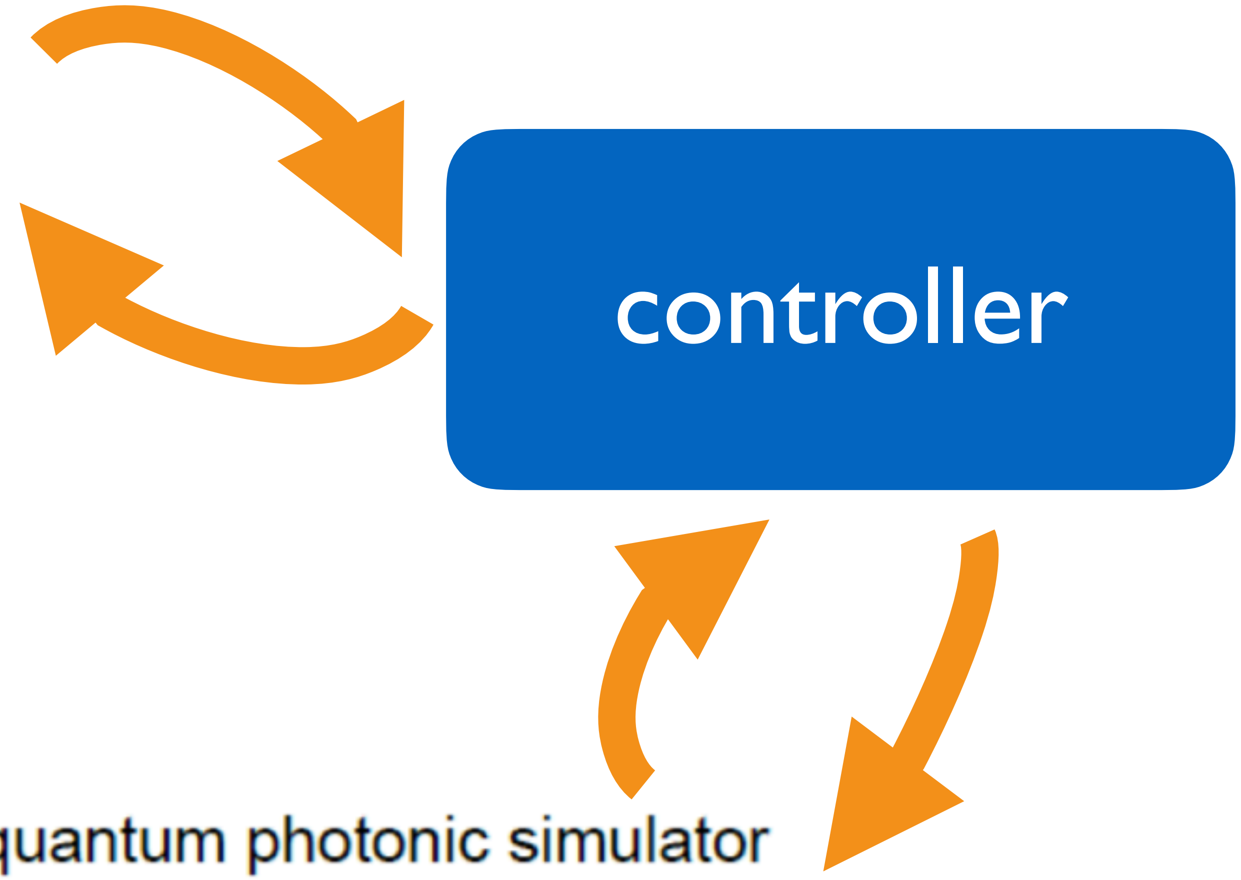

Si quantum photonic simulator

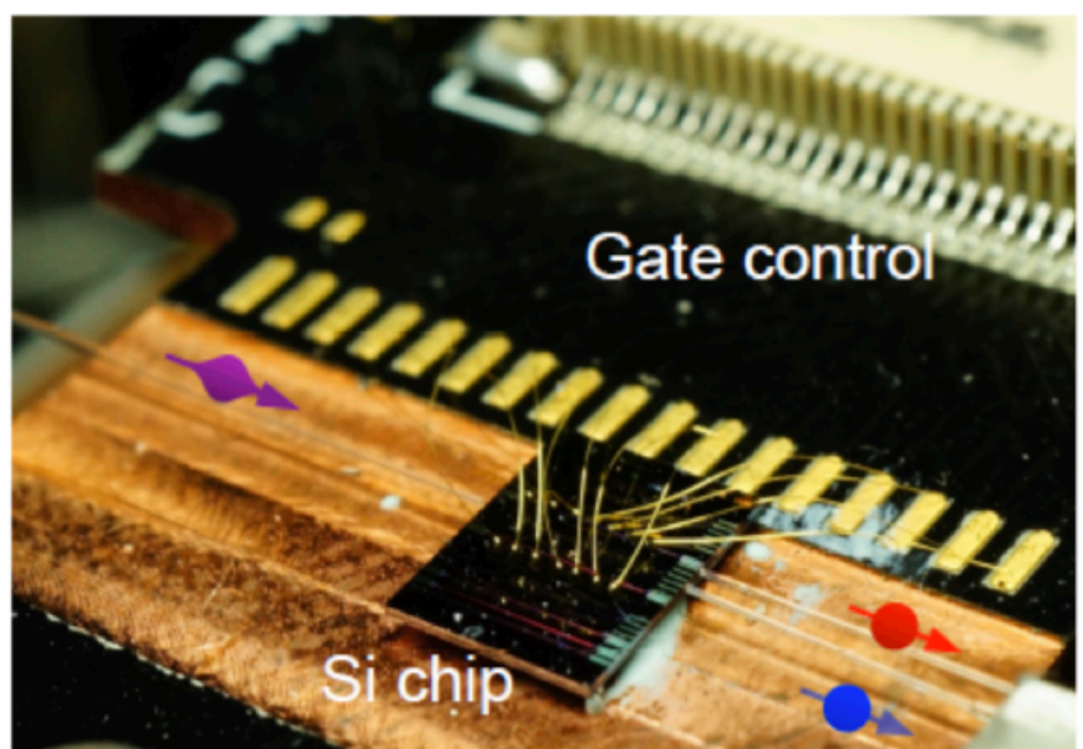

Wang, Paesani, Santagati et al (Thompson group Bristol) 2017

Producing new experimental layouts

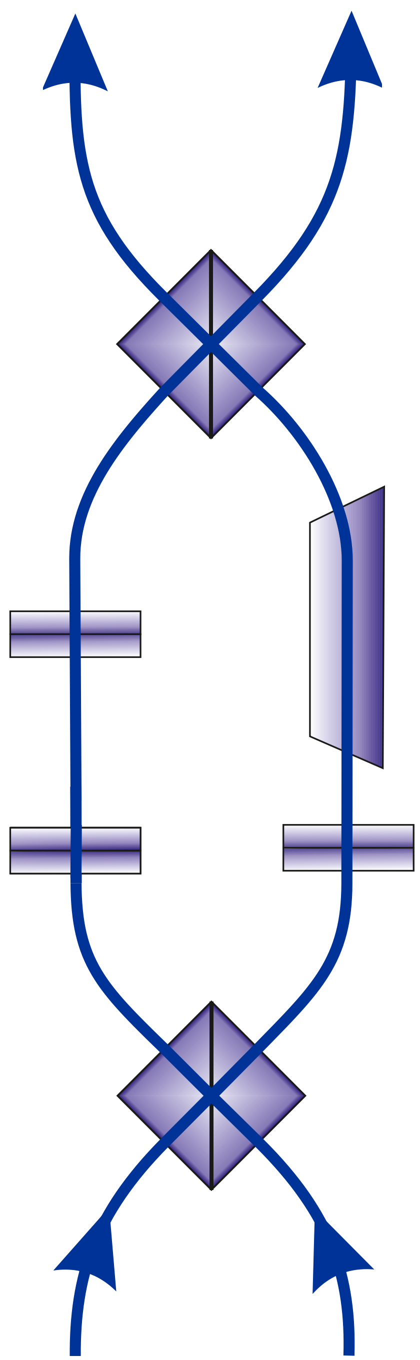

combine optical elements to produce highly entangled states

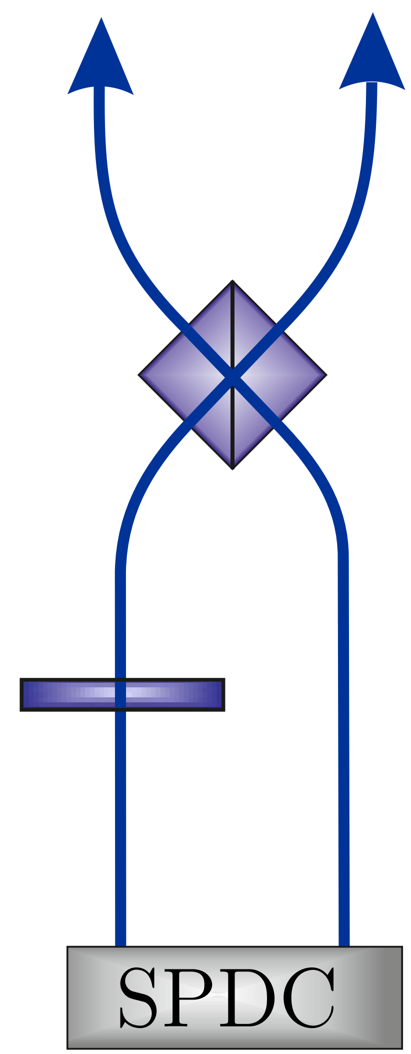

Krenn et al 2016

Briegel group

Melnikov et al PNAS 2018

Quantum control and feedback

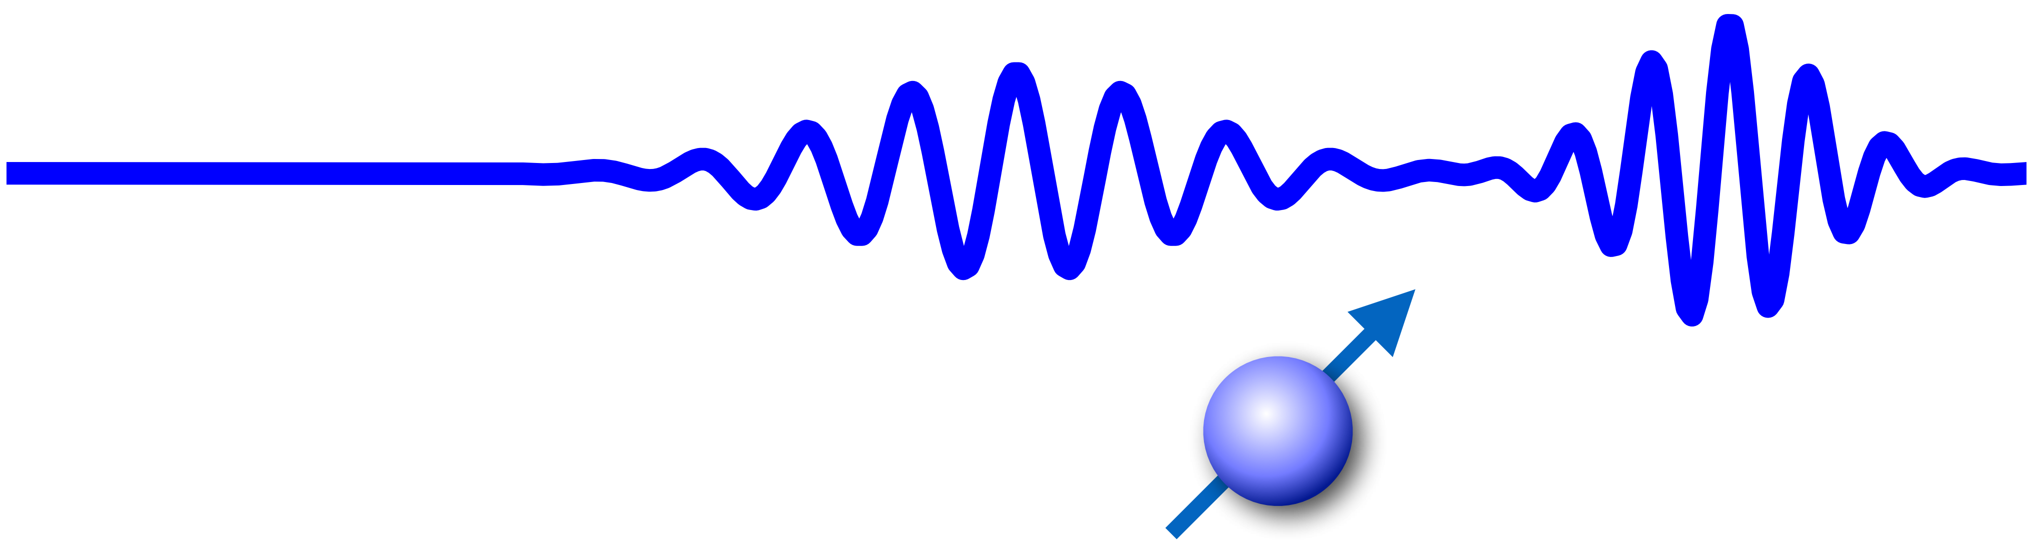

open-loop
(no feedback)

closed-loop
(with feedback)

typical tasks: state preparation, unitary synthesis, state/subspace stabilization, feedback cooling and initialization, quantum error correction

Quantum control and feedback

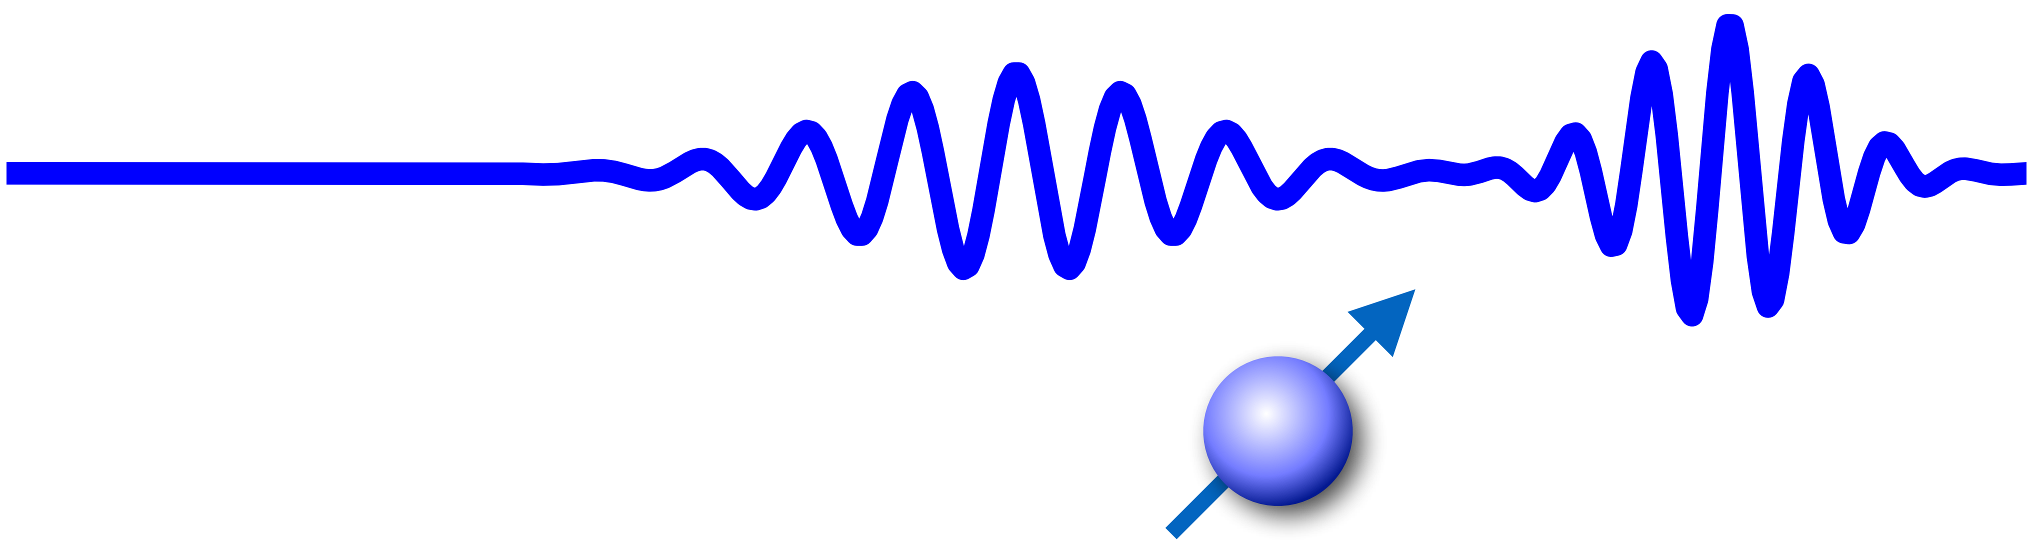

traditional: numerical techniques like GRAPE

new machine-learning techniques:

- model-free (implicitly learn model from behaviour)
- can easily include feedback
- profit from computer science method development

Quantum control

Qubit control

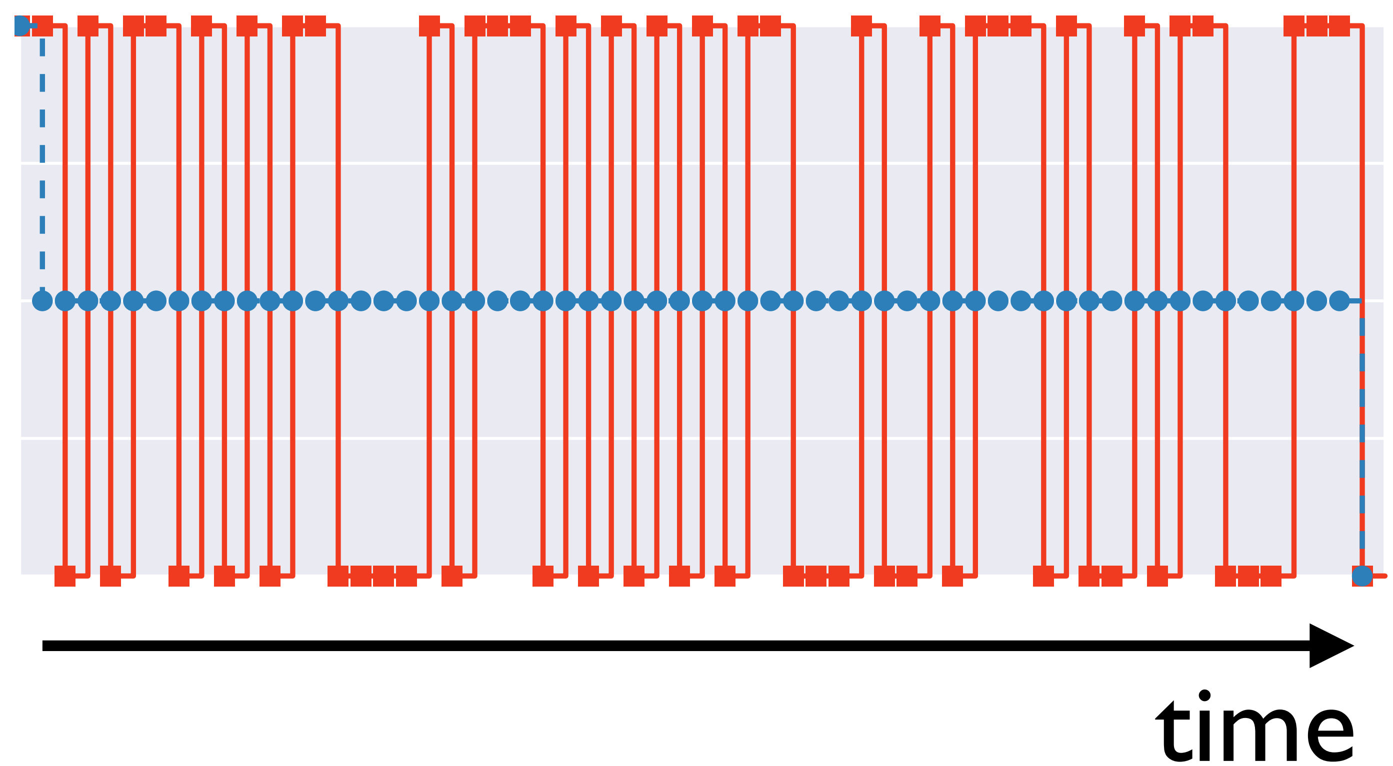

Bukov et al PRX 2018

August, Hernandez-Lobato 2018

Niu et al 2019

...

Quantum control

one aspect of 'quantum machine learning':
optimizing control parameters
directly on a quantum device

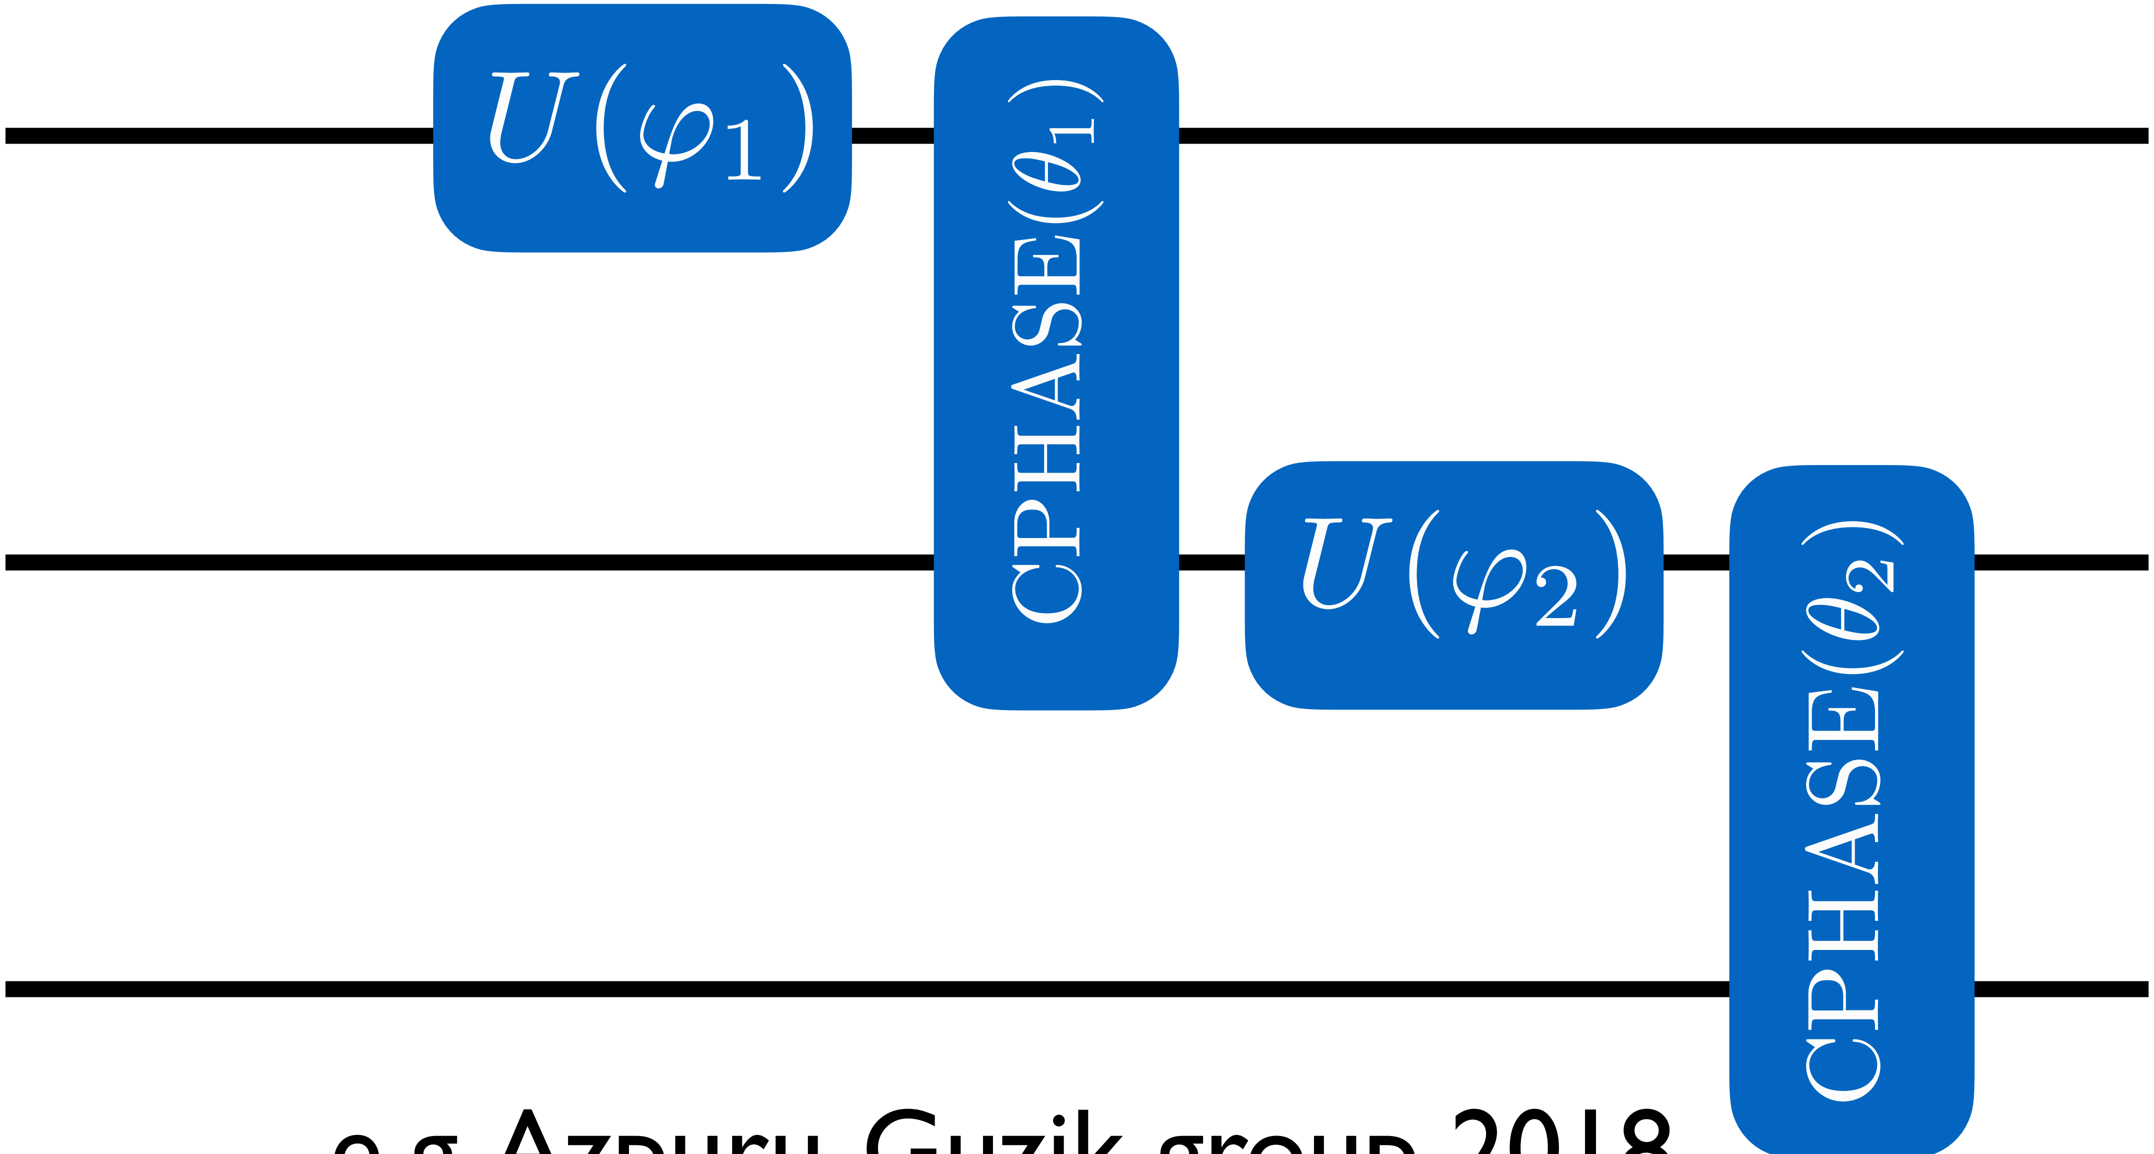

e.g. Azpuru-Guzik group 2018

Deep Reinforcement Learning

“Supervised learning”

(most neural network applications)

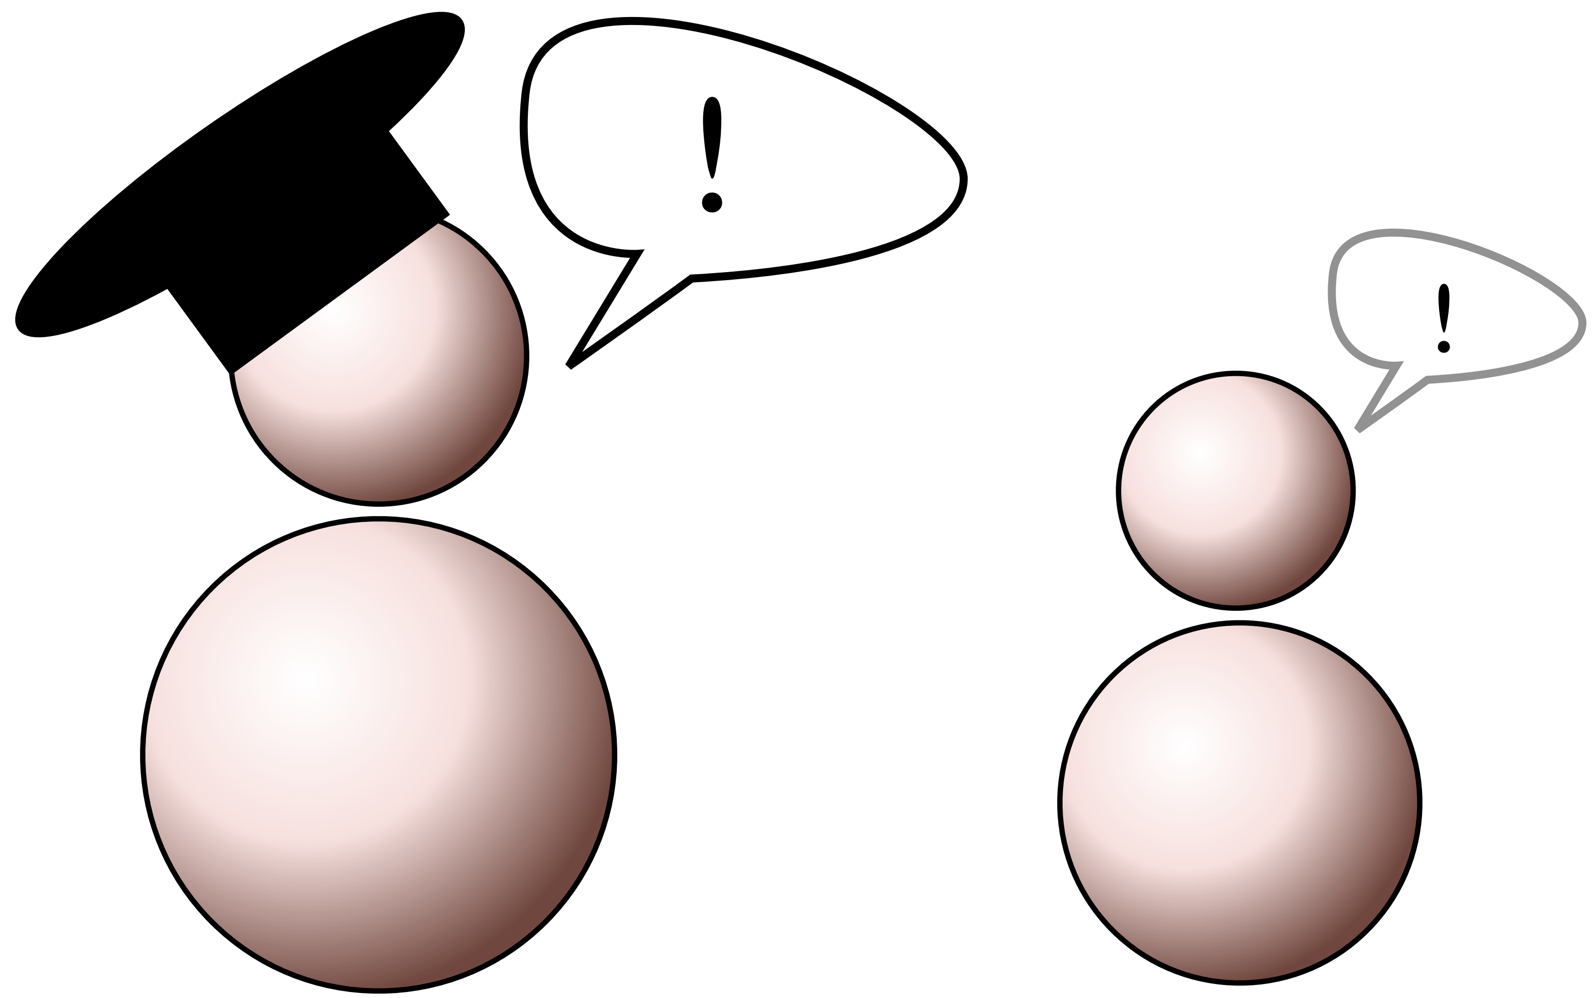

teacher
(smart)

student
(imitates teacher)

final level limited by teacher

“Reinforcement learning”

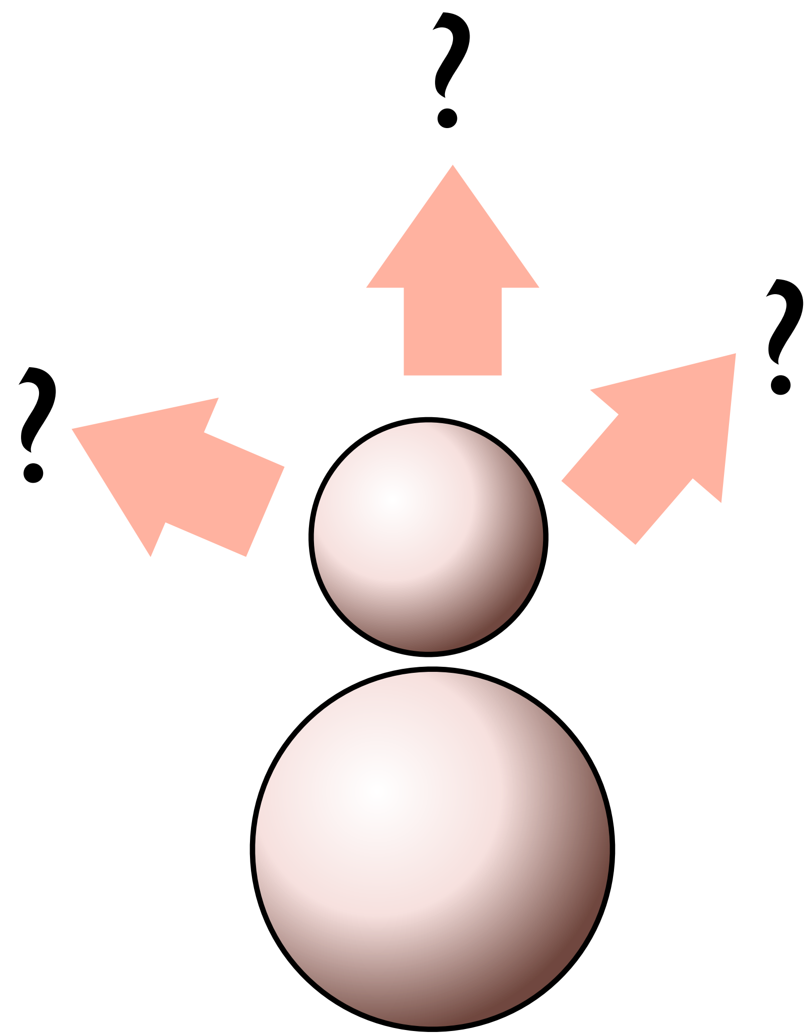

student/scientist
(tries out things)

final level: unlimited (?)


Image: Wikipedia
AlphaGo 2017

observe state, pick action, get reward

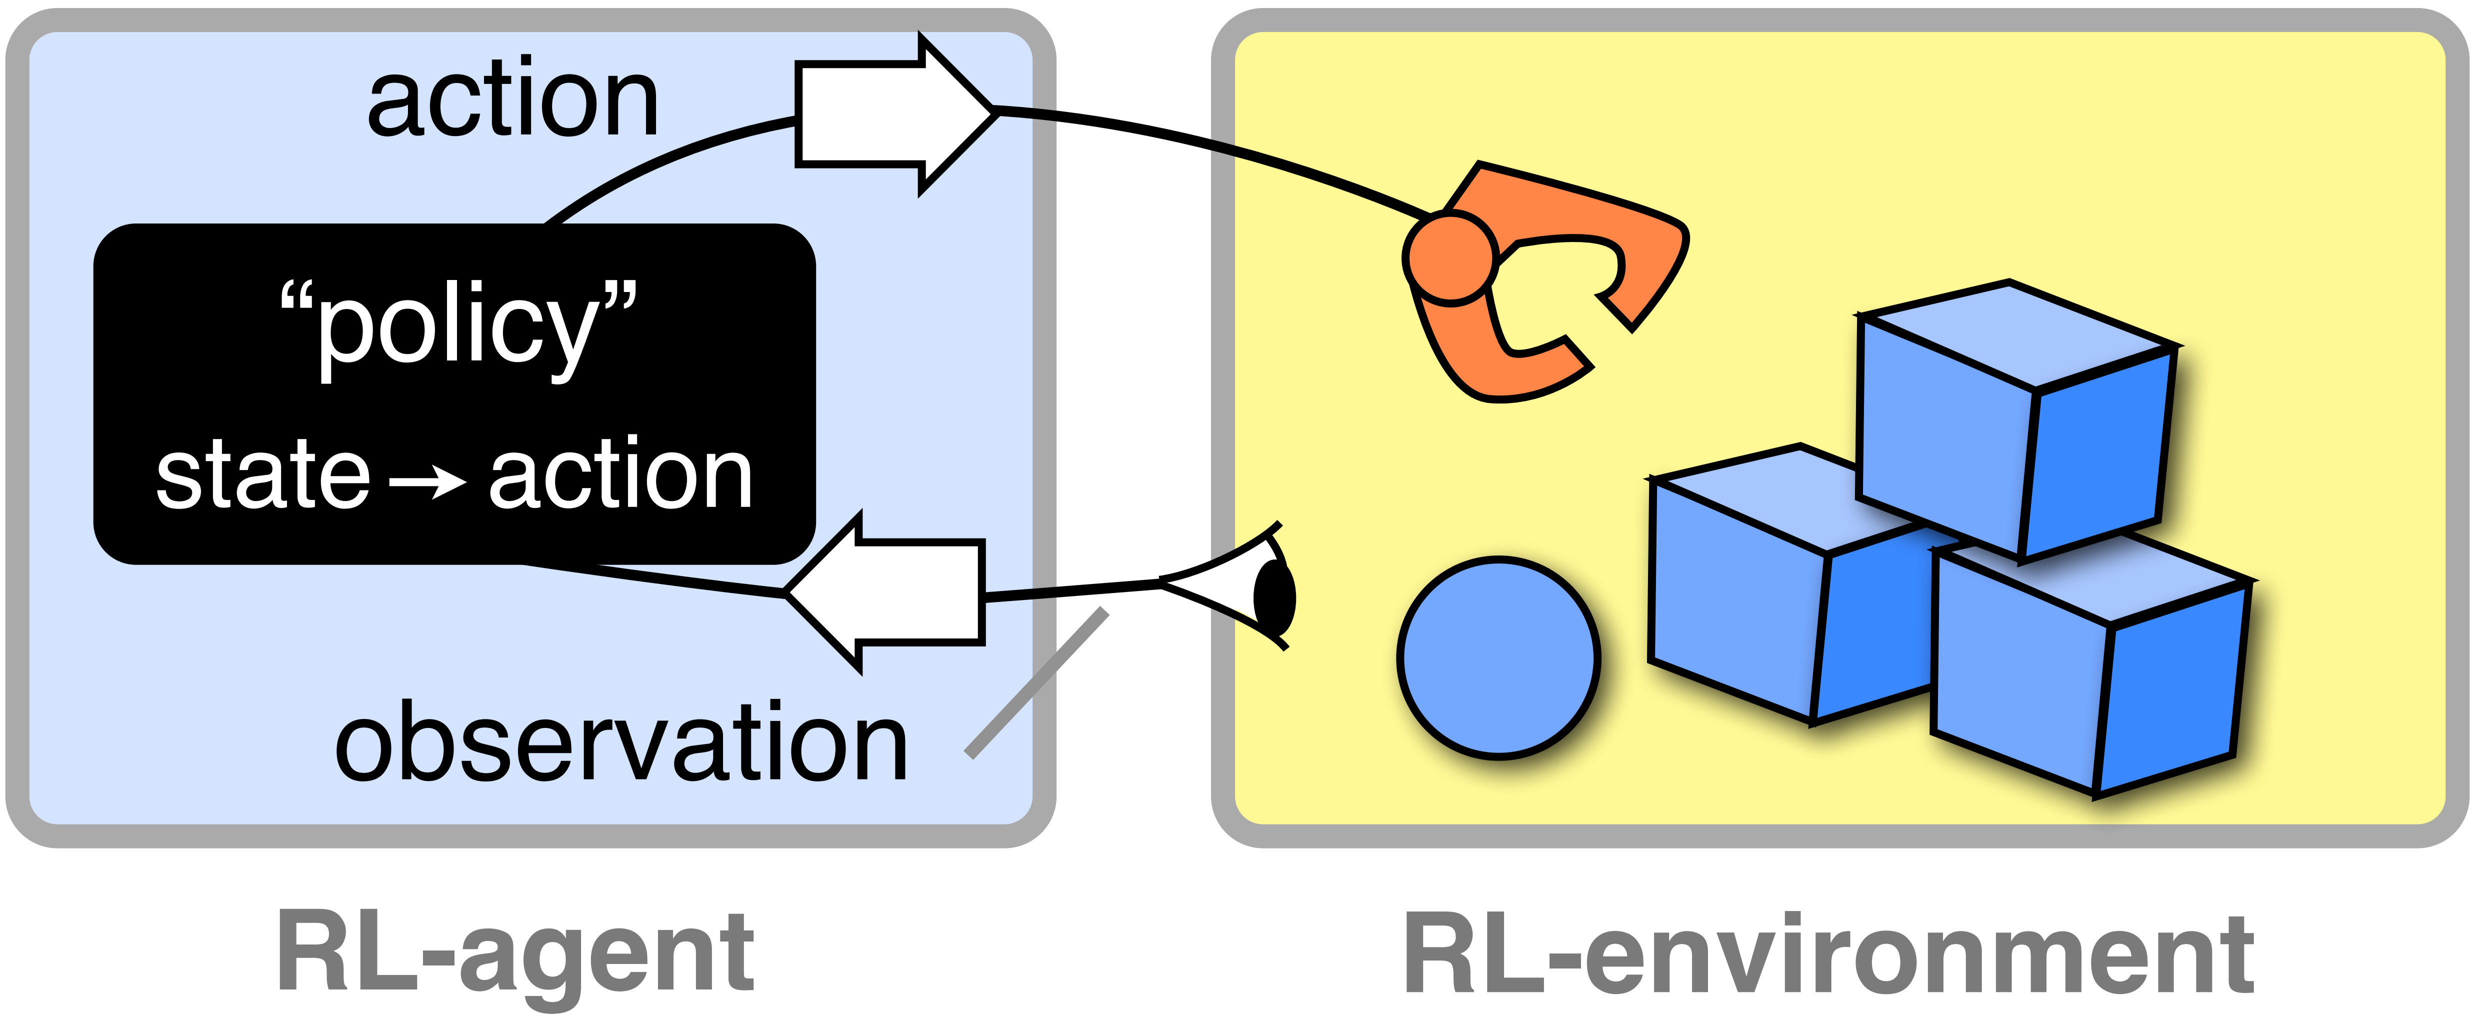

Policy: $\pi_\theta(a_t|s_t)$ – probability to pick action a_t given observed state s_t at time t

Maximize expected "return" R : sum of rewards

$$\delta\theta_j = \frac{\partial}{\partial\theta_j}\bar{R} = \sum_t \mathbb{E}\left[R\frac{\partial}{\partial\theta_j}\ln\pi_\theta(a_t|s_t)\right]$$

"Policy gradient", "REINFORCE" (Williams 1992)

Today: **deep** reinforcement learning
with deep neural networks: high-dimensional
states and actions

Reinforcement learning: Advantages

Discover **Feedback** Strategies

(beyond GRAPE etc.)

No feedback: A^N strategies (A #actions N #steps)

With feedback: A^{M^N} strategies (M #msmt outcomes)

Model-free

No need to develop/fit/calibrate model/equations for dynamics of the world/the device

…can learn on real devices, with all imperfections

Reinforcement learning: Advantages

RL with deep neural networks:
Handle **arbitrary observations**

(images, videos, measurement results of any kind, sentences, graphs, quantum states, quantum circuits…)

Reinforcement learning: Challenges

Need to see **many evolutions**!
tens of thousands

Cannot discover 'isolated/rare-event' strategies
(also true for any other non-domain-specific algorithm)

Discovering Quantum Error Correction Strategies

work with **Thomas Fösel, Petru Tighineanu, and Talitha Weiss**

Physical Review X 031084 (2018)

Quantum Error Correction: many approaches

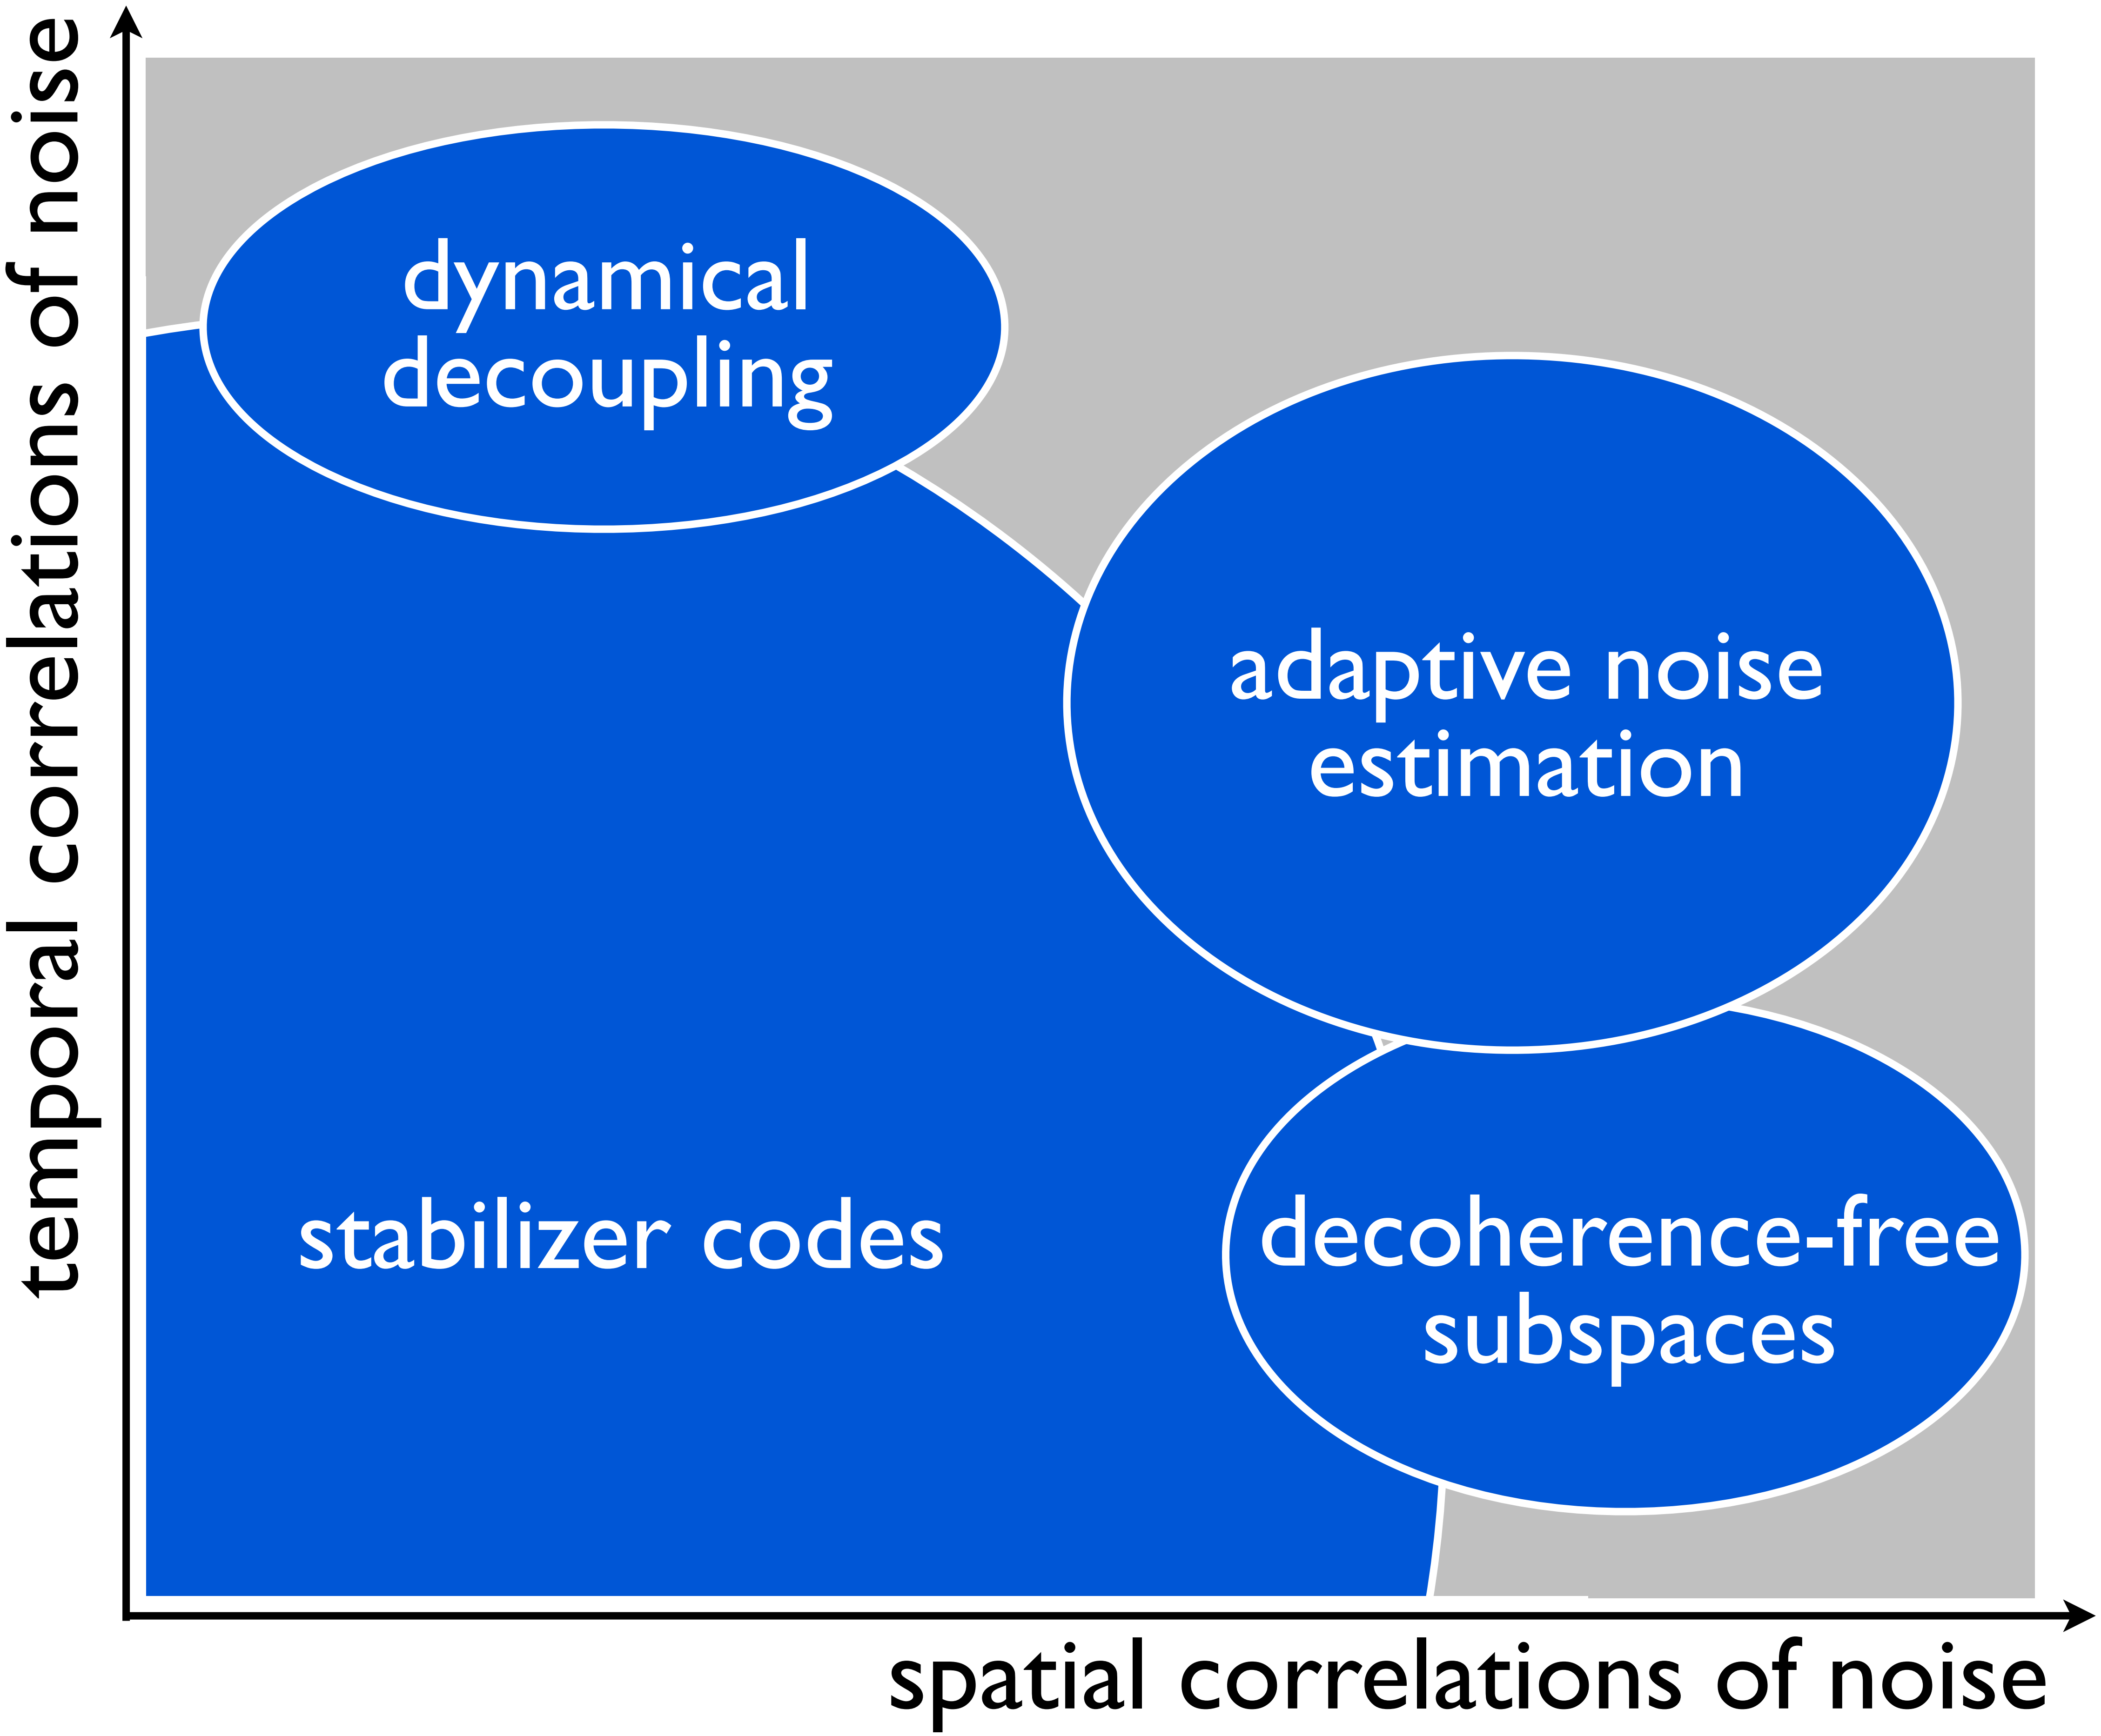

Quantum Error Correction: many approaches

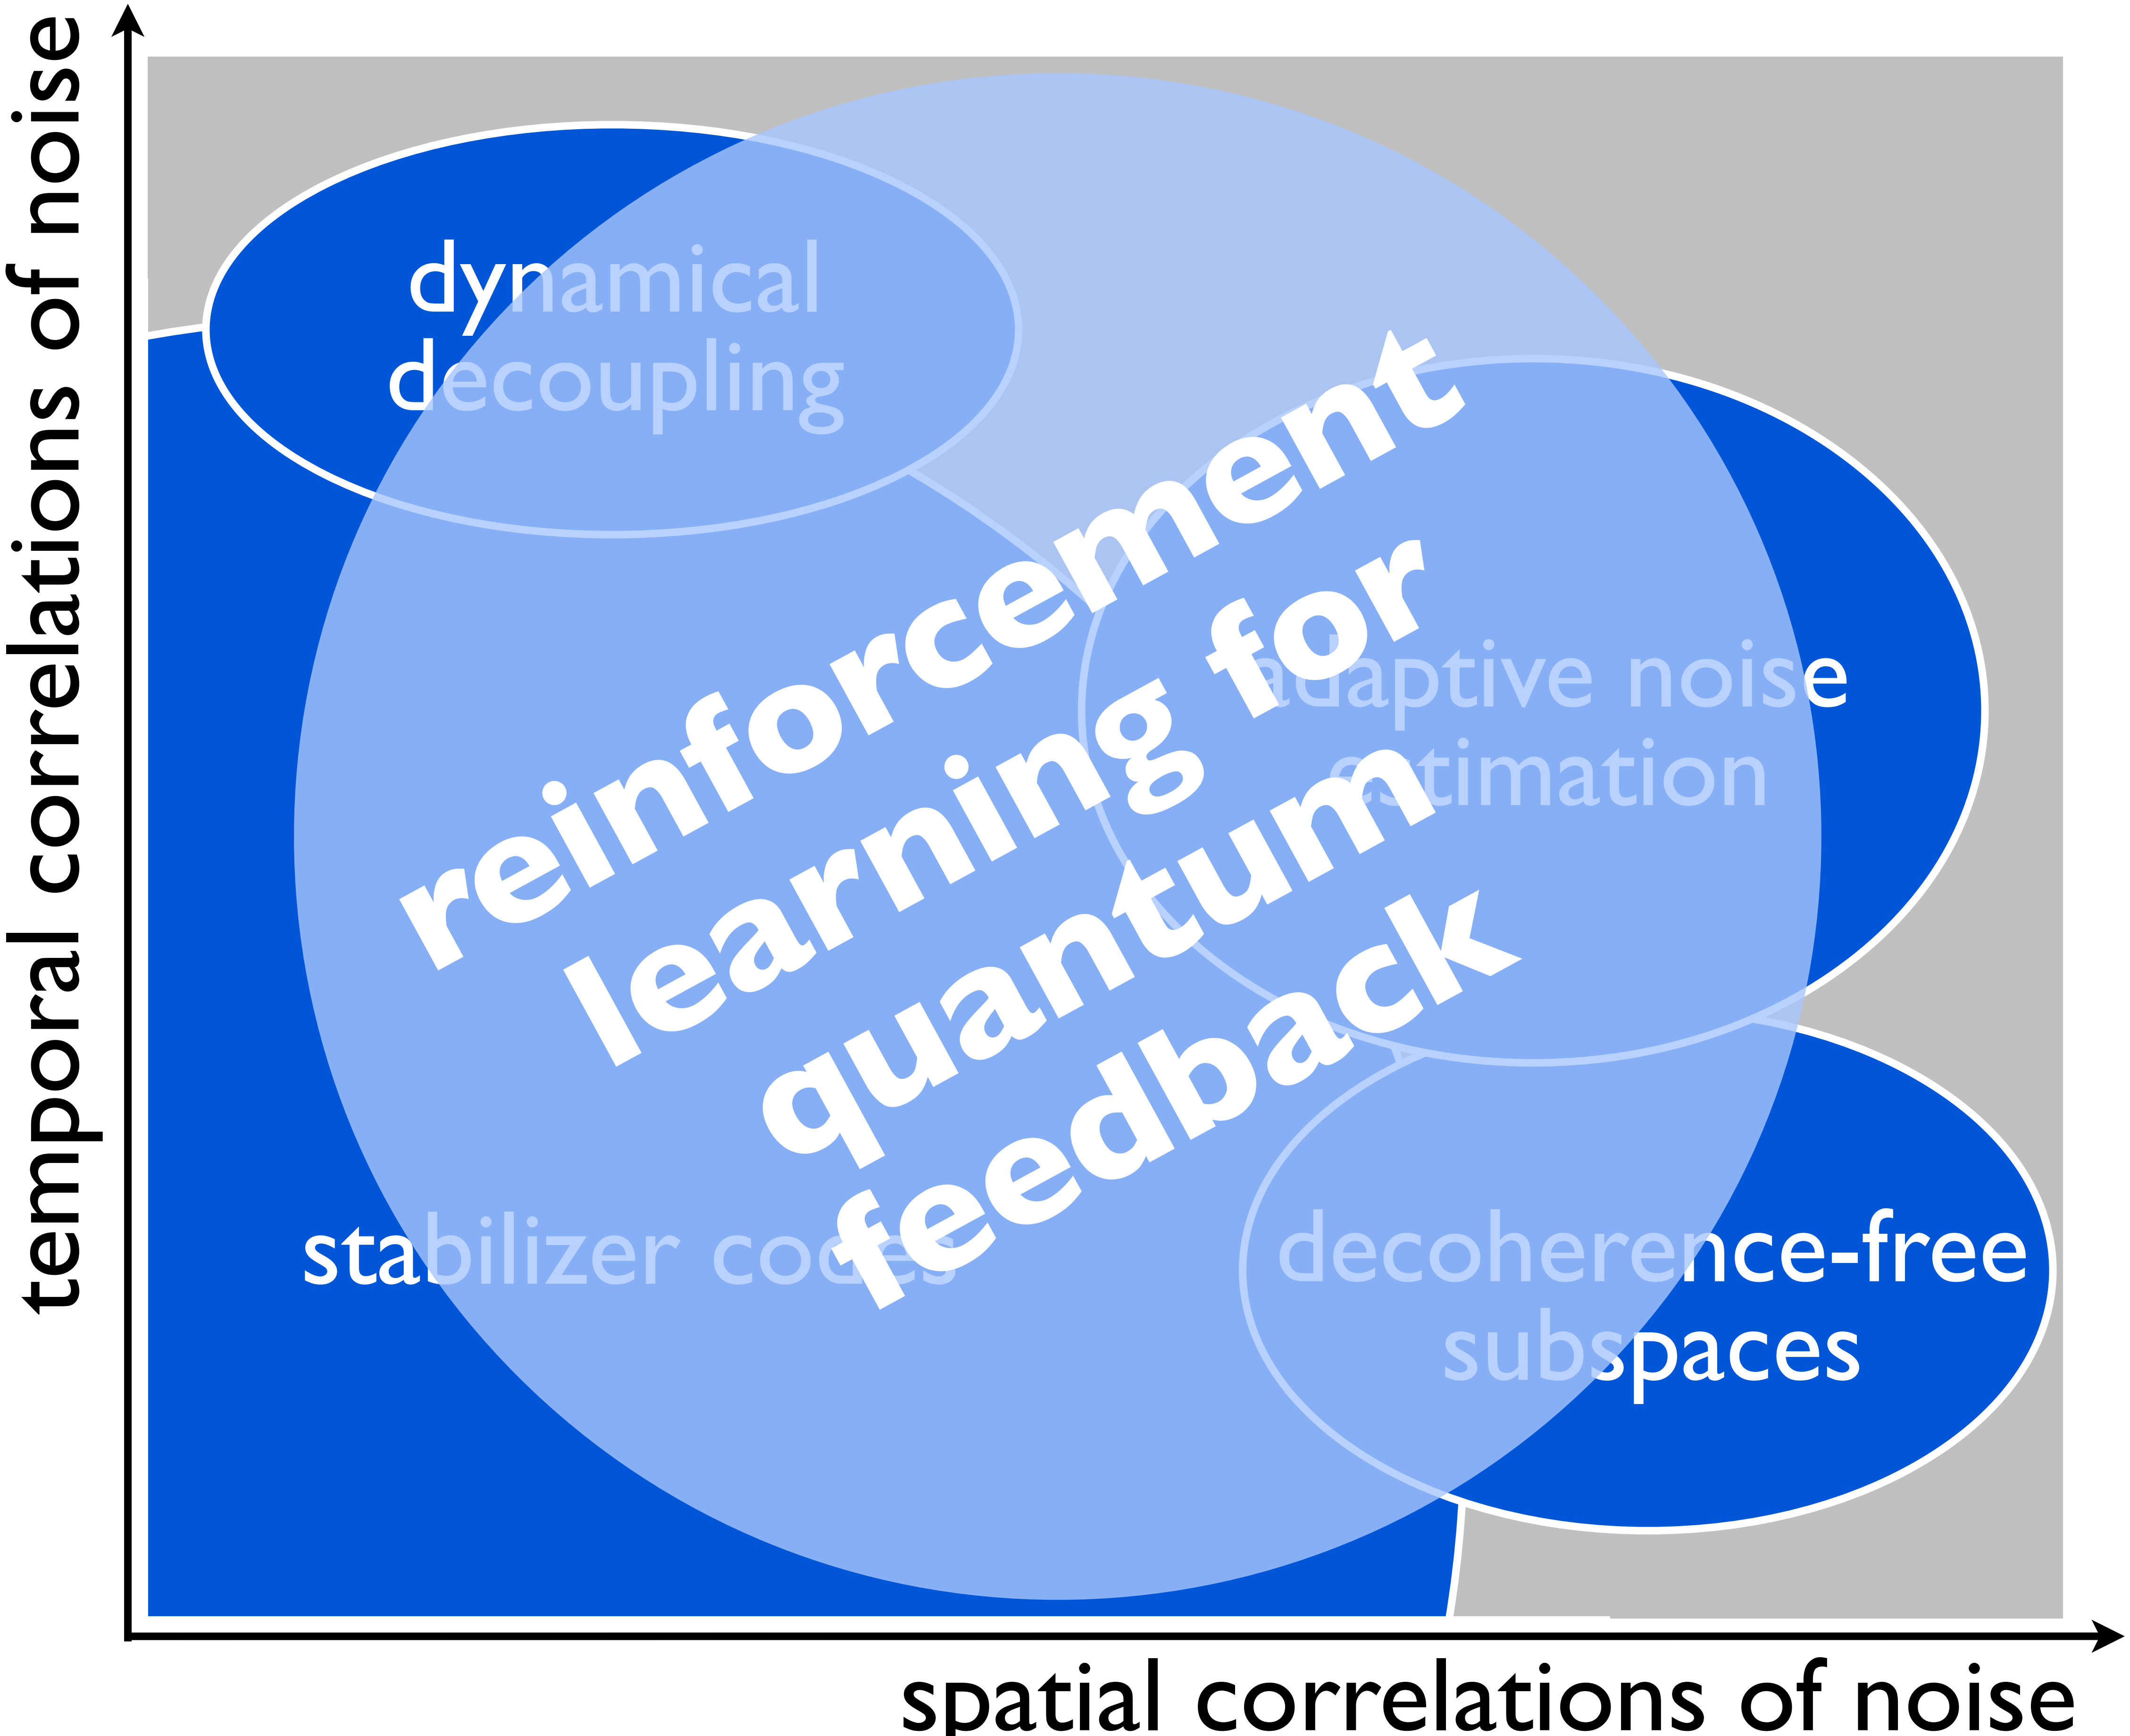

Goal: protect a small quantum module against noise

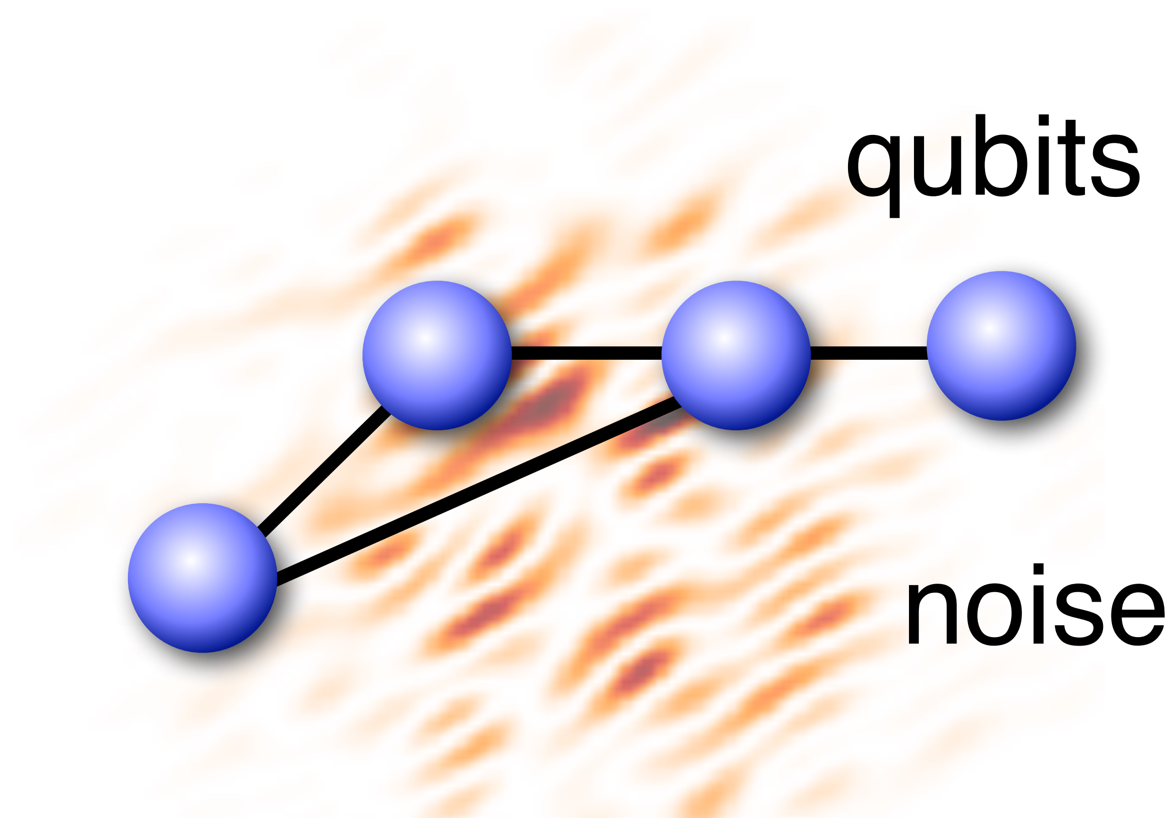

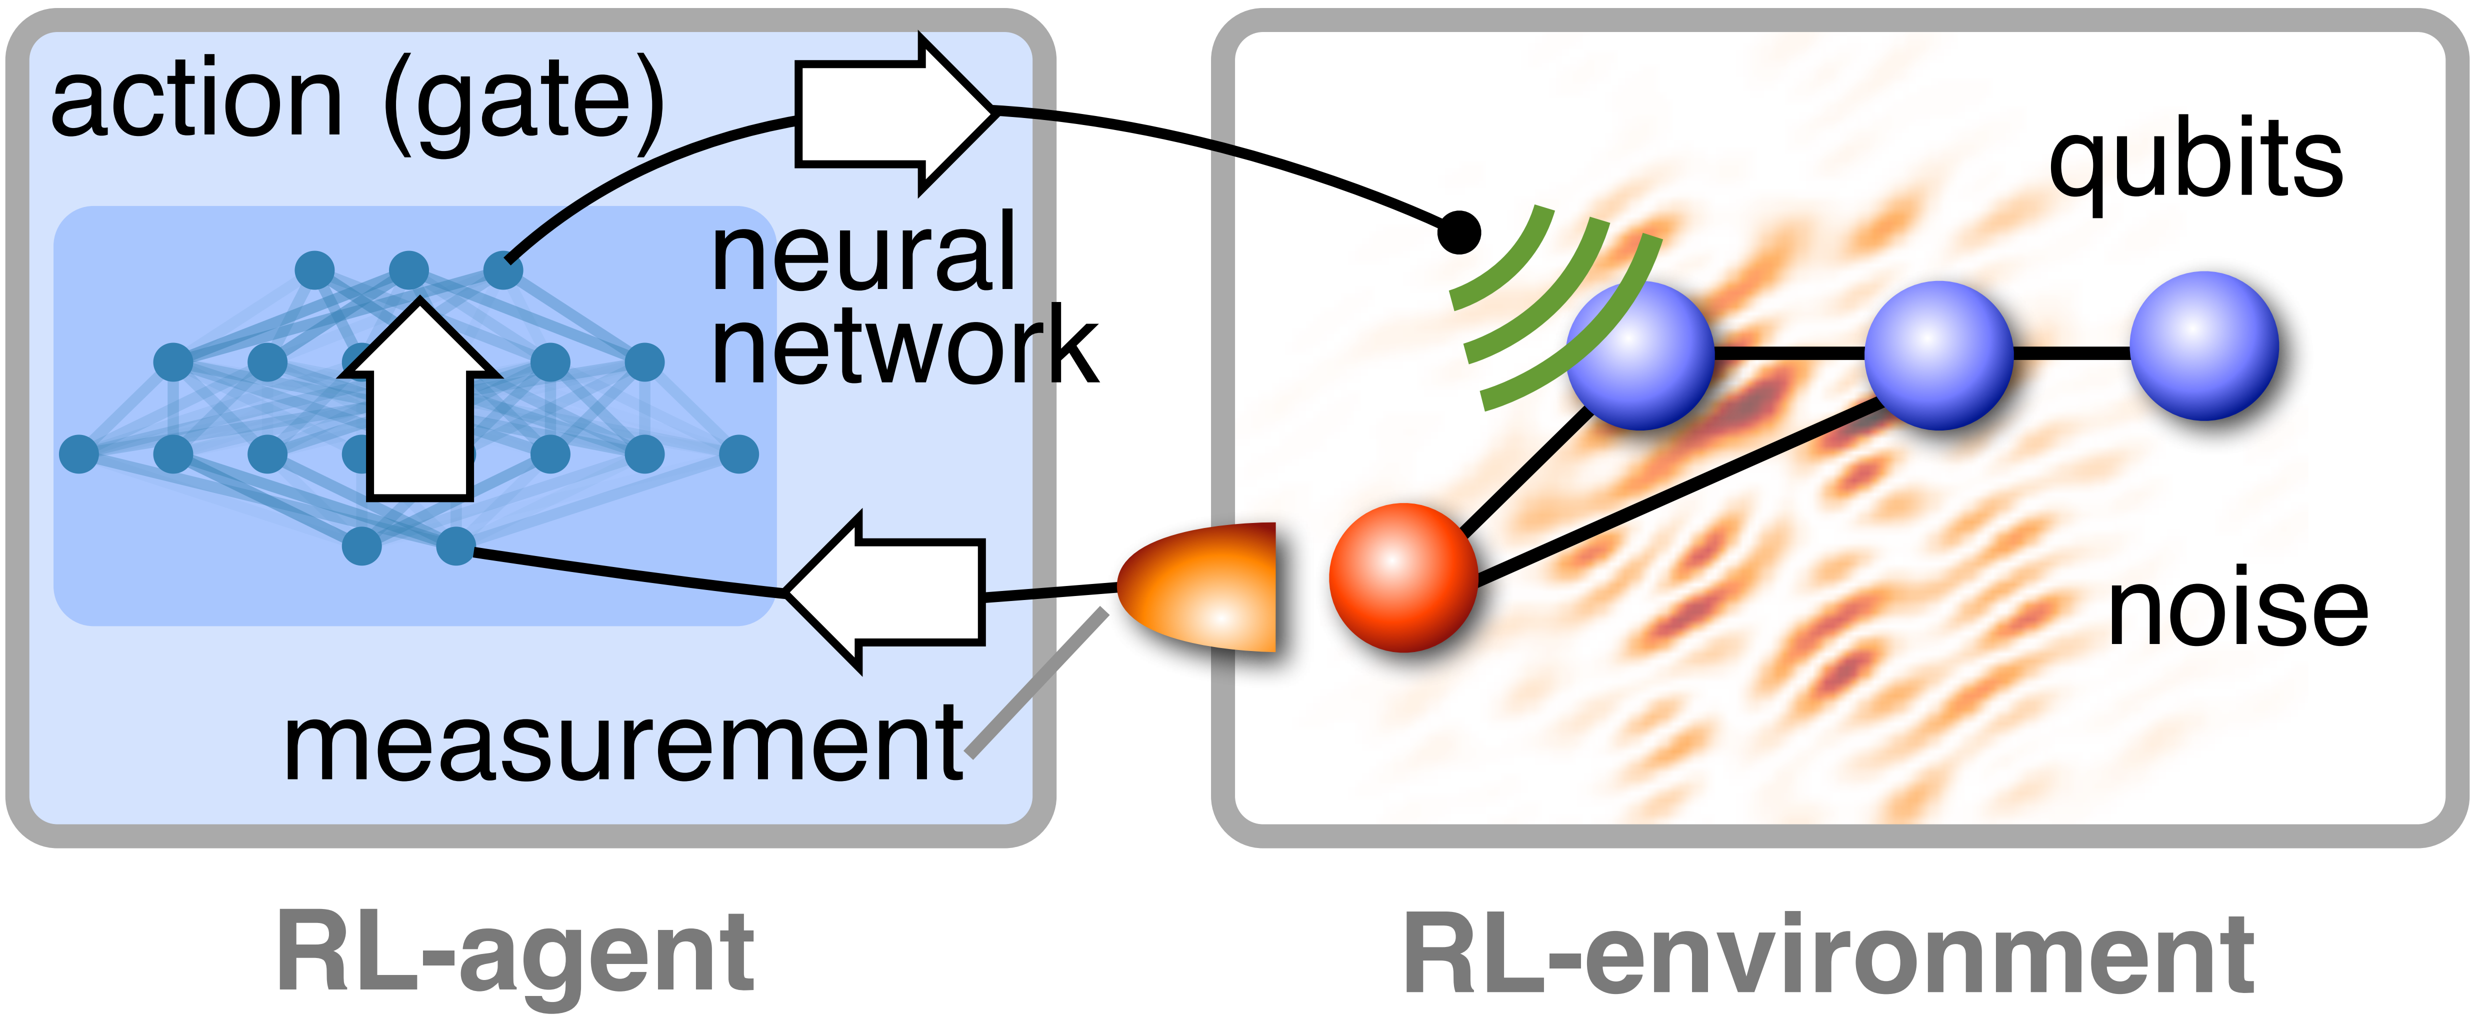

try out actions, get reward, improve strategy, …

Initialize qubit here: $|\Psi(0)\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$

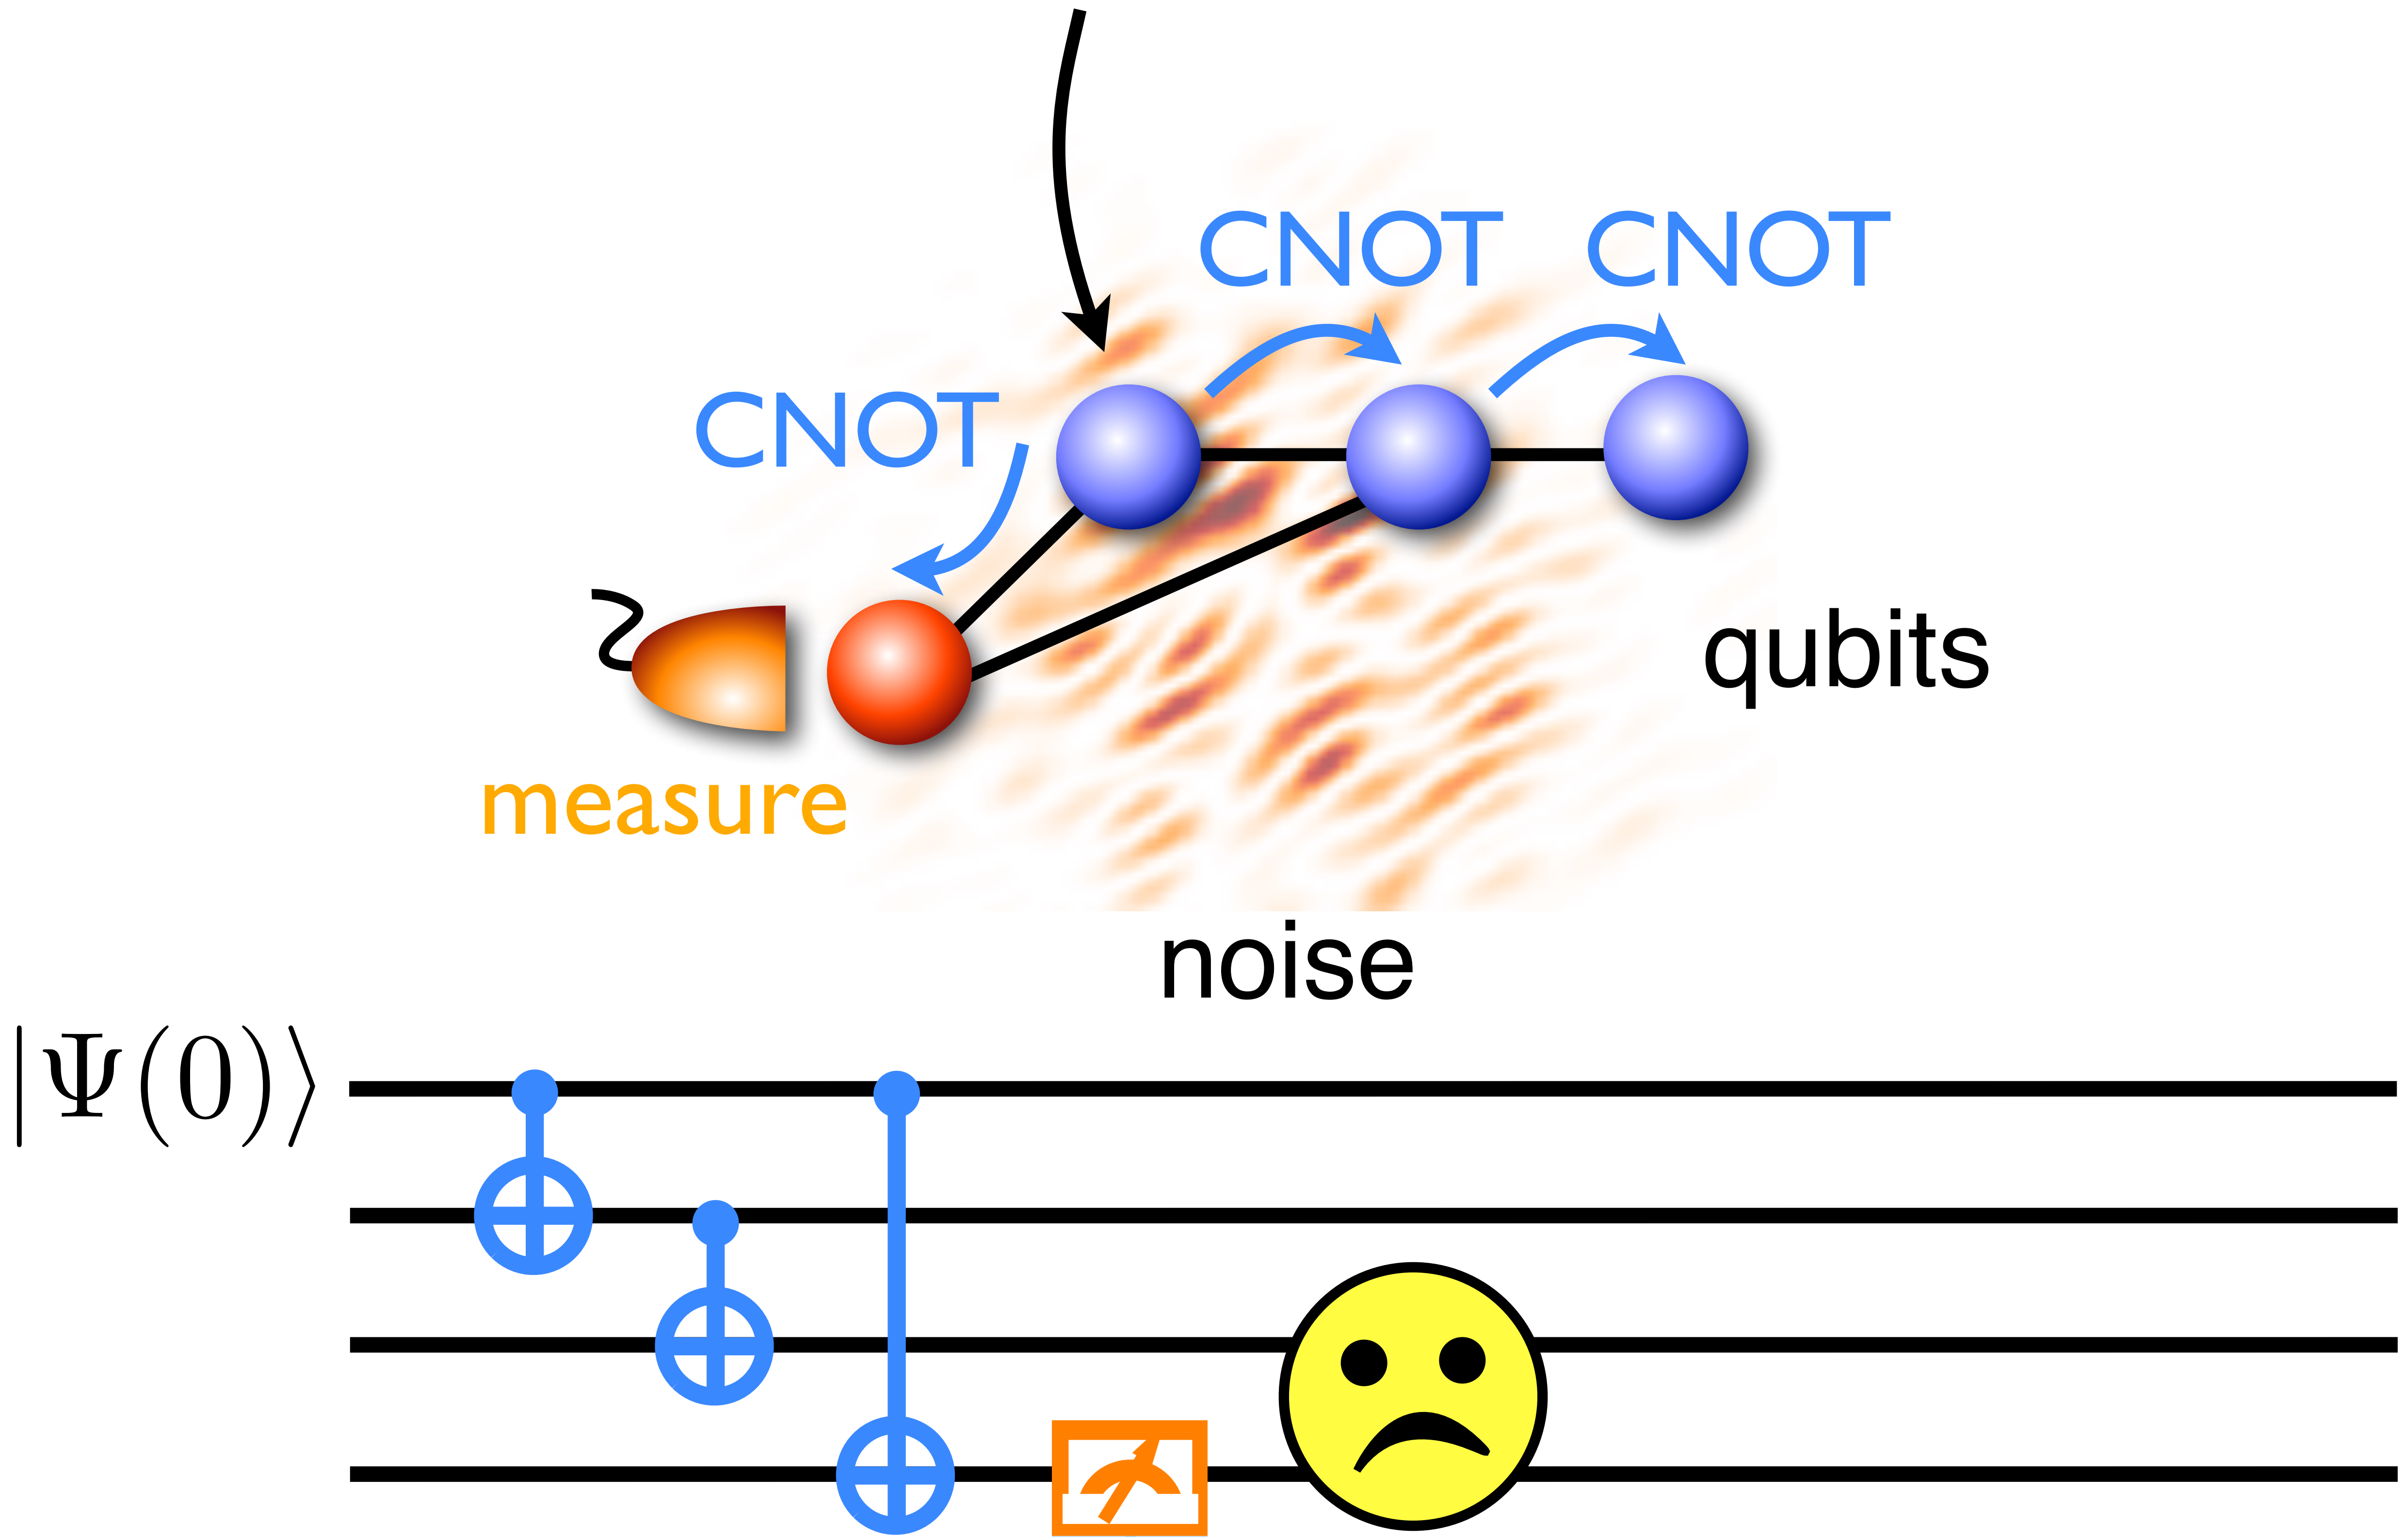

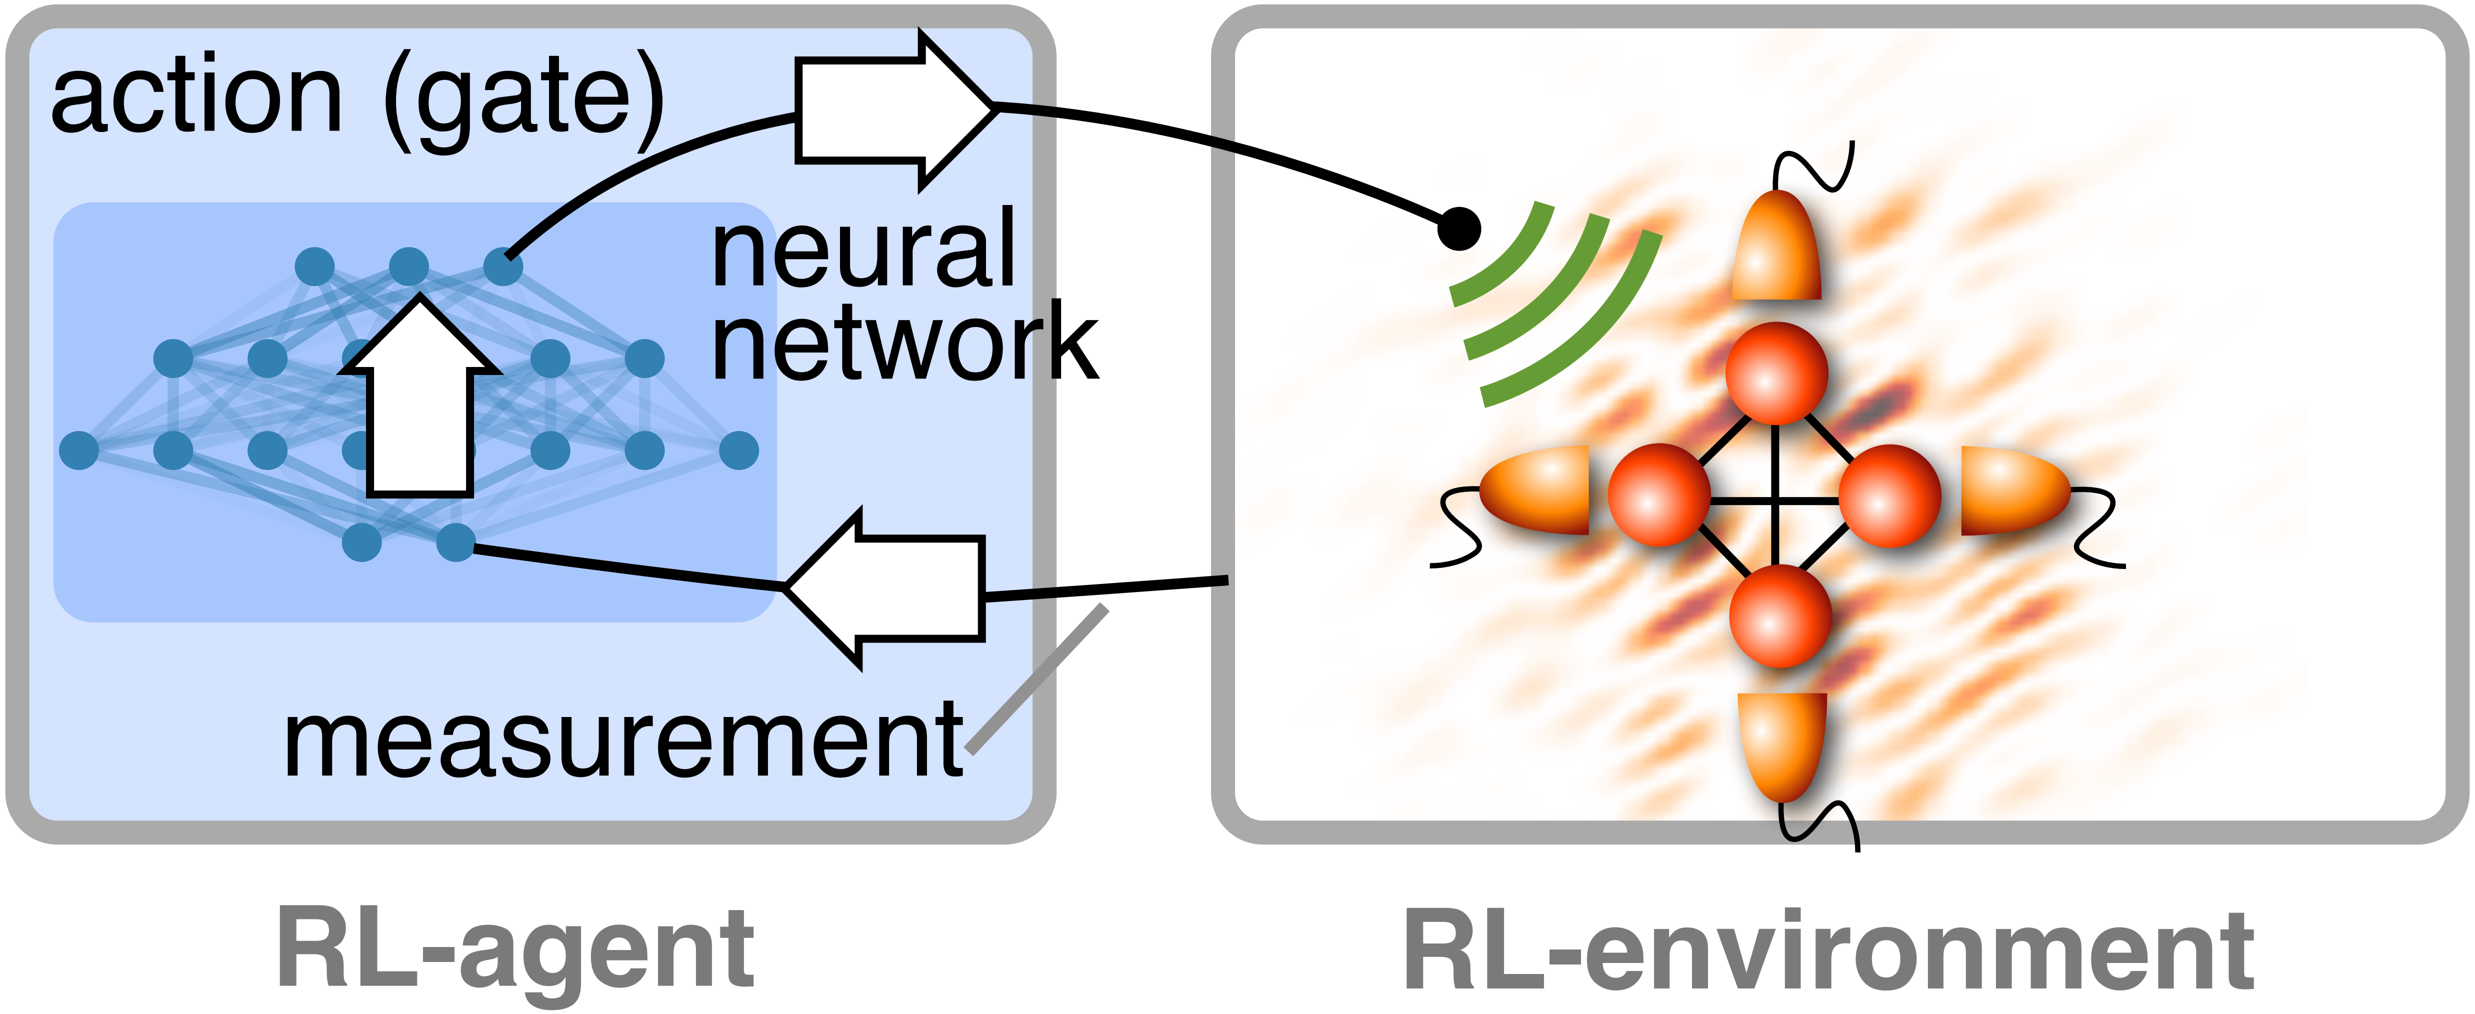

example: 4 qubits, measurements possible on all, CNOTs between all, bit-flip noise on all

$$\dot{\hat{\rho}} = \frac{1}{T_{\mathrm{dec}}} \sum_j (\hat{\sigma}_{xj} \hat{\rho} \hat{\sigma}_{xj} - \hat{\rho})$$

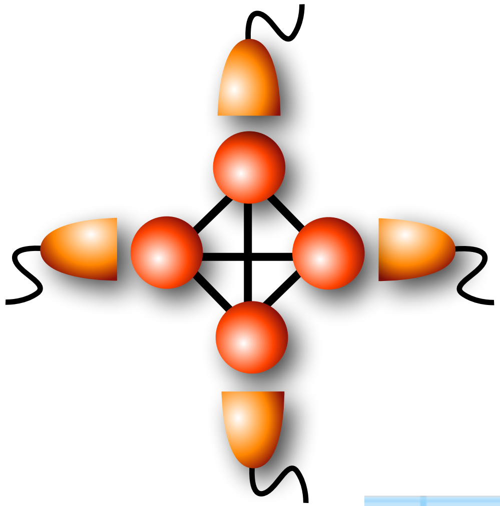

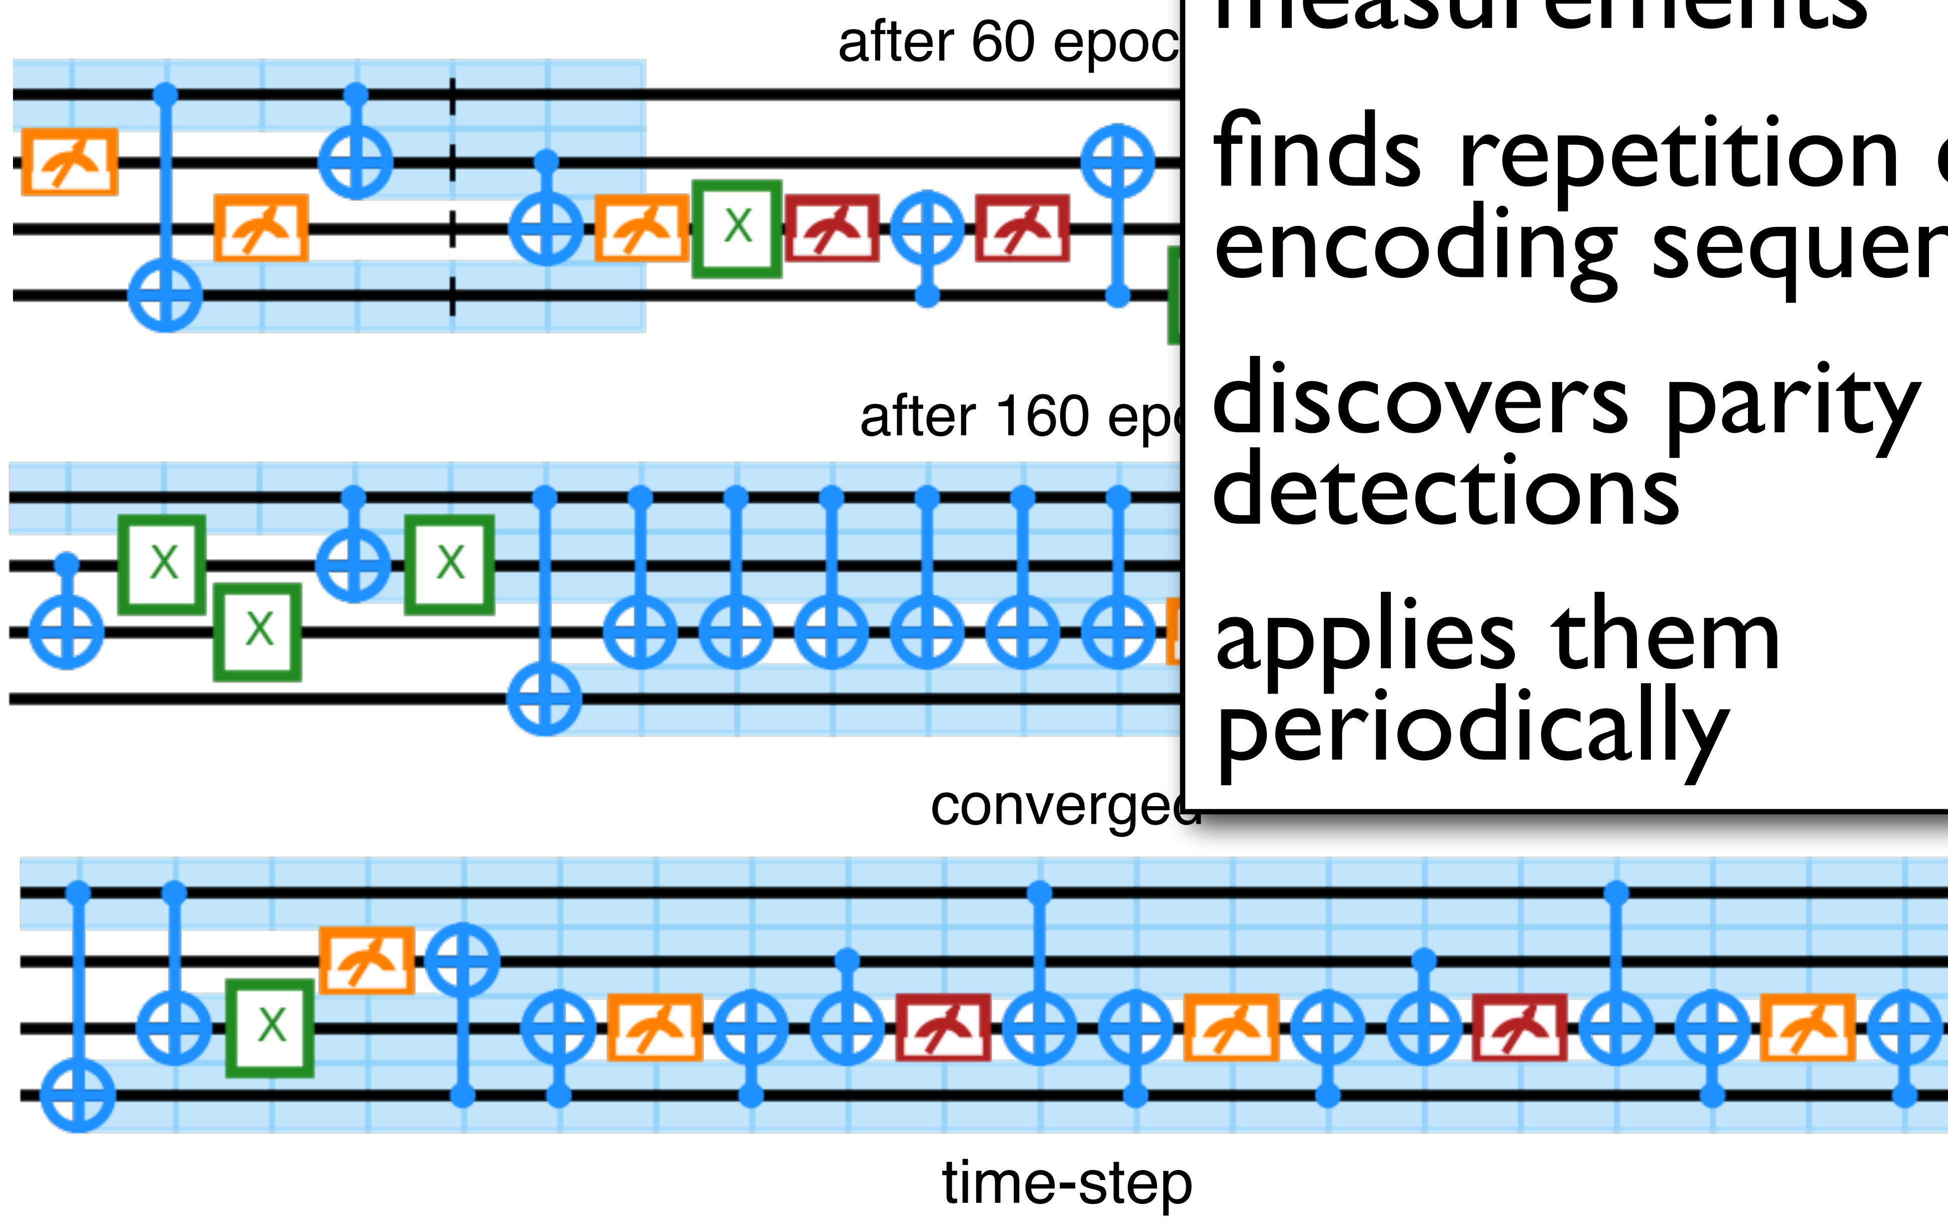

The network...

avoids catastrophic measurements

finds repetition code encoding sequence

discovers parity detections

applies them periodically

(natural policy gradient)

Physical Review X 031084 (2018)

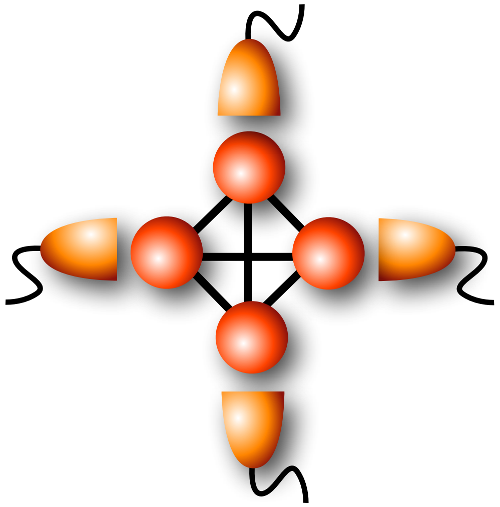

Training Progress

measure of success=
"recoverable quantum information"

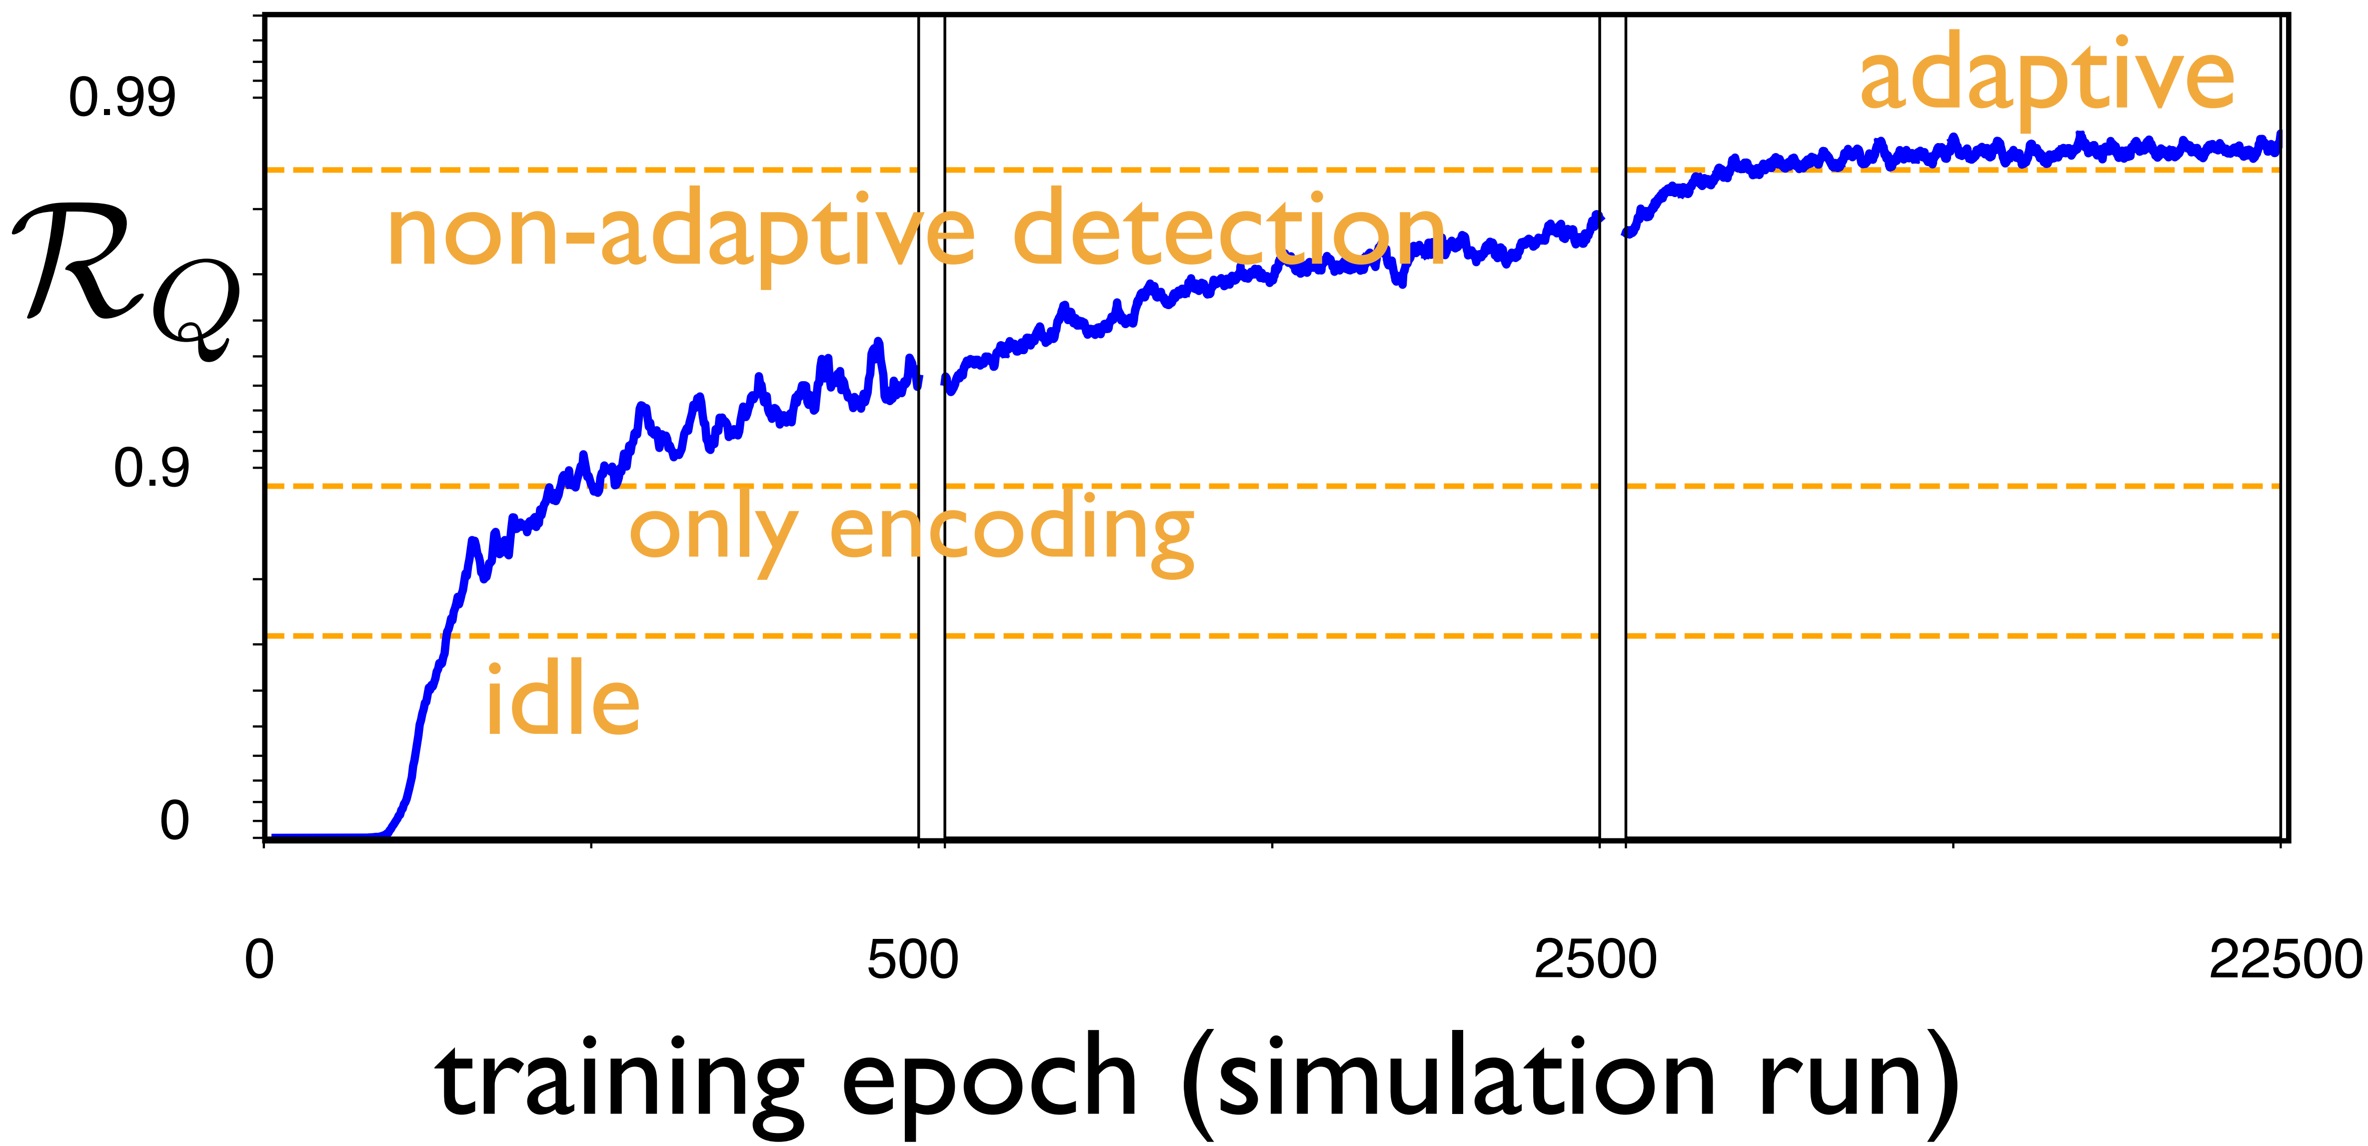

Physical Review X 031084 (2018)

Different topologies

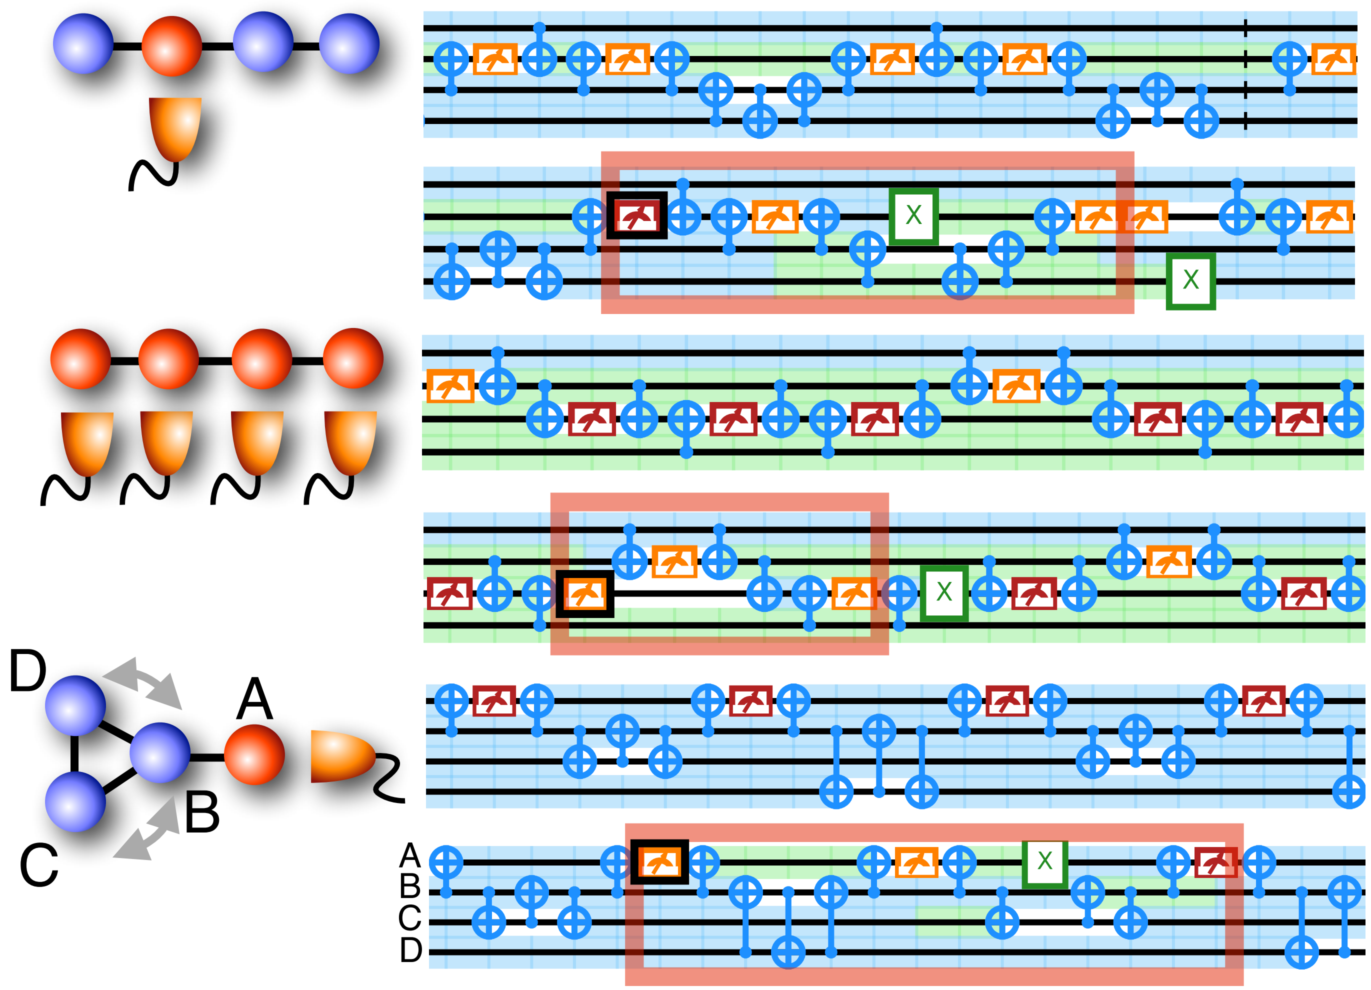

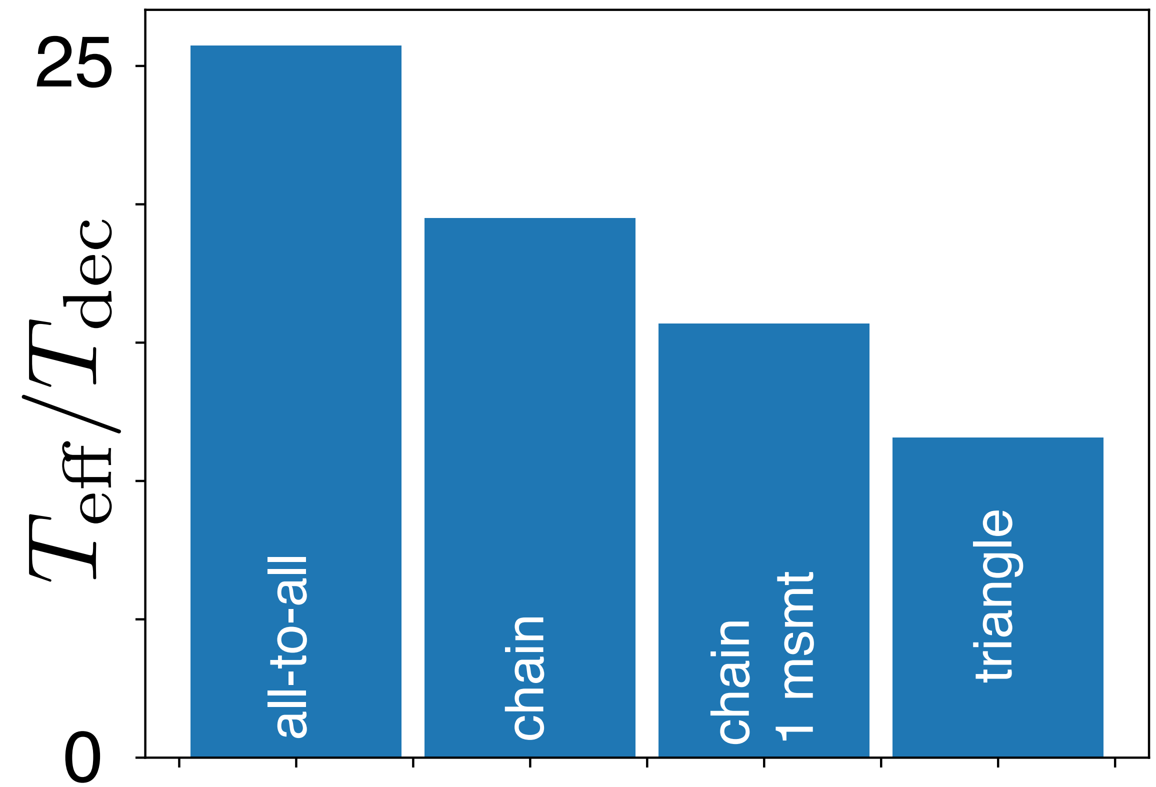

Physical Review X 031084 (2018)

Different class of scenarios: Dephasing by a noisy field

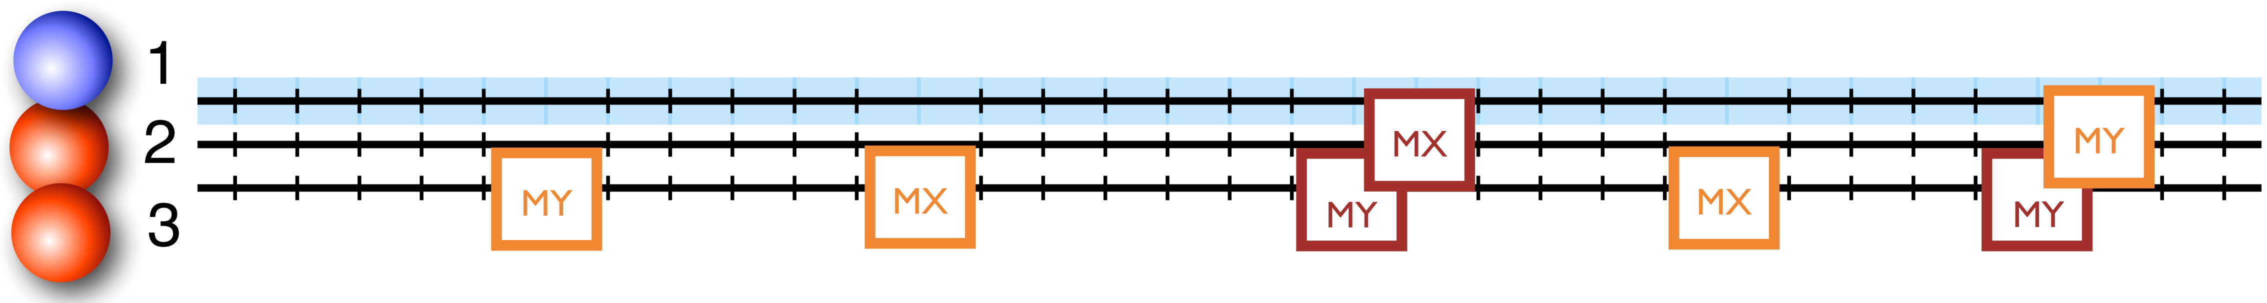

Network discovers **adaptive** noise estimation strategy

Strategy = decision tree

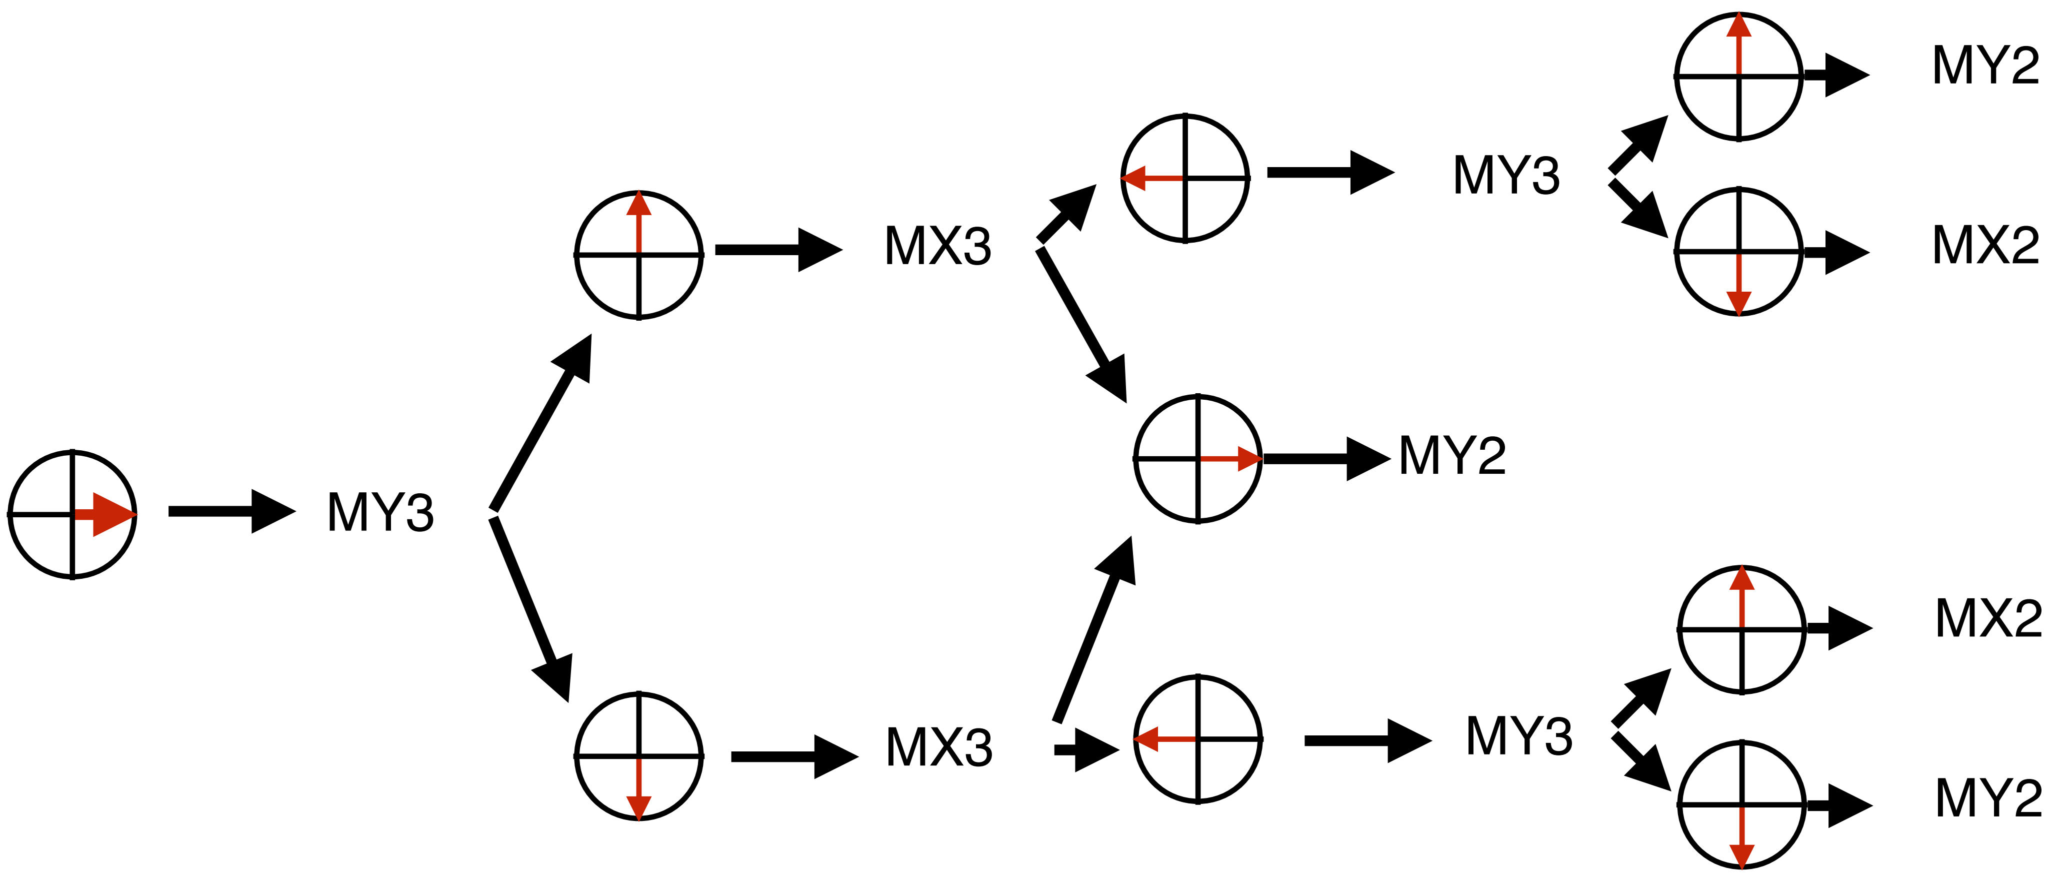

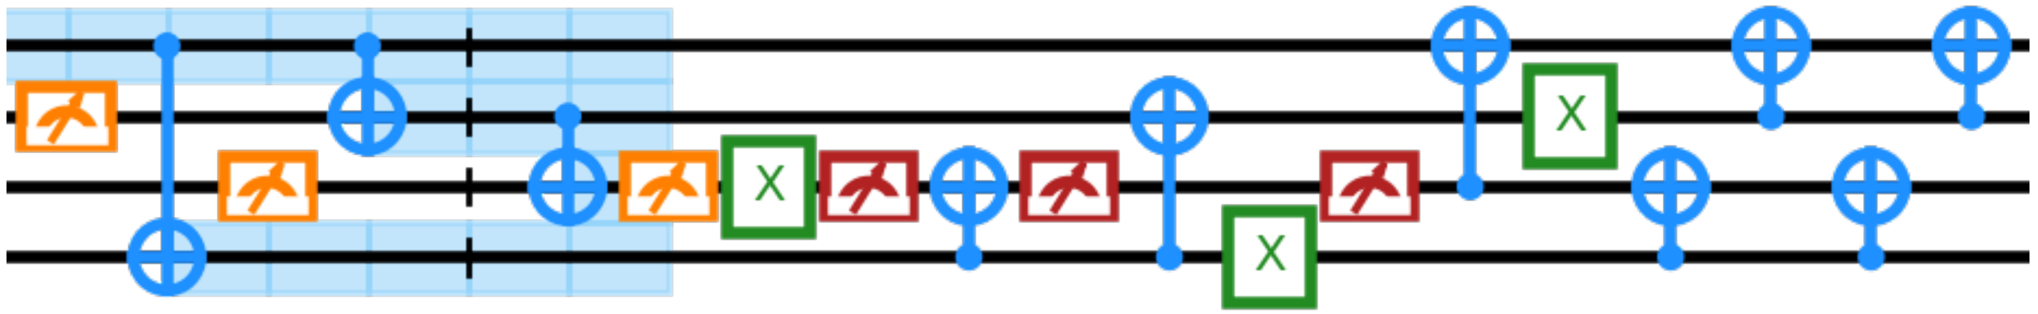

Two key concepts to make it work

Construct smart reward
("Recoverable Quantum Information")

As much information as possible &
two-stage learning

Physical Review X 031084 (2018)

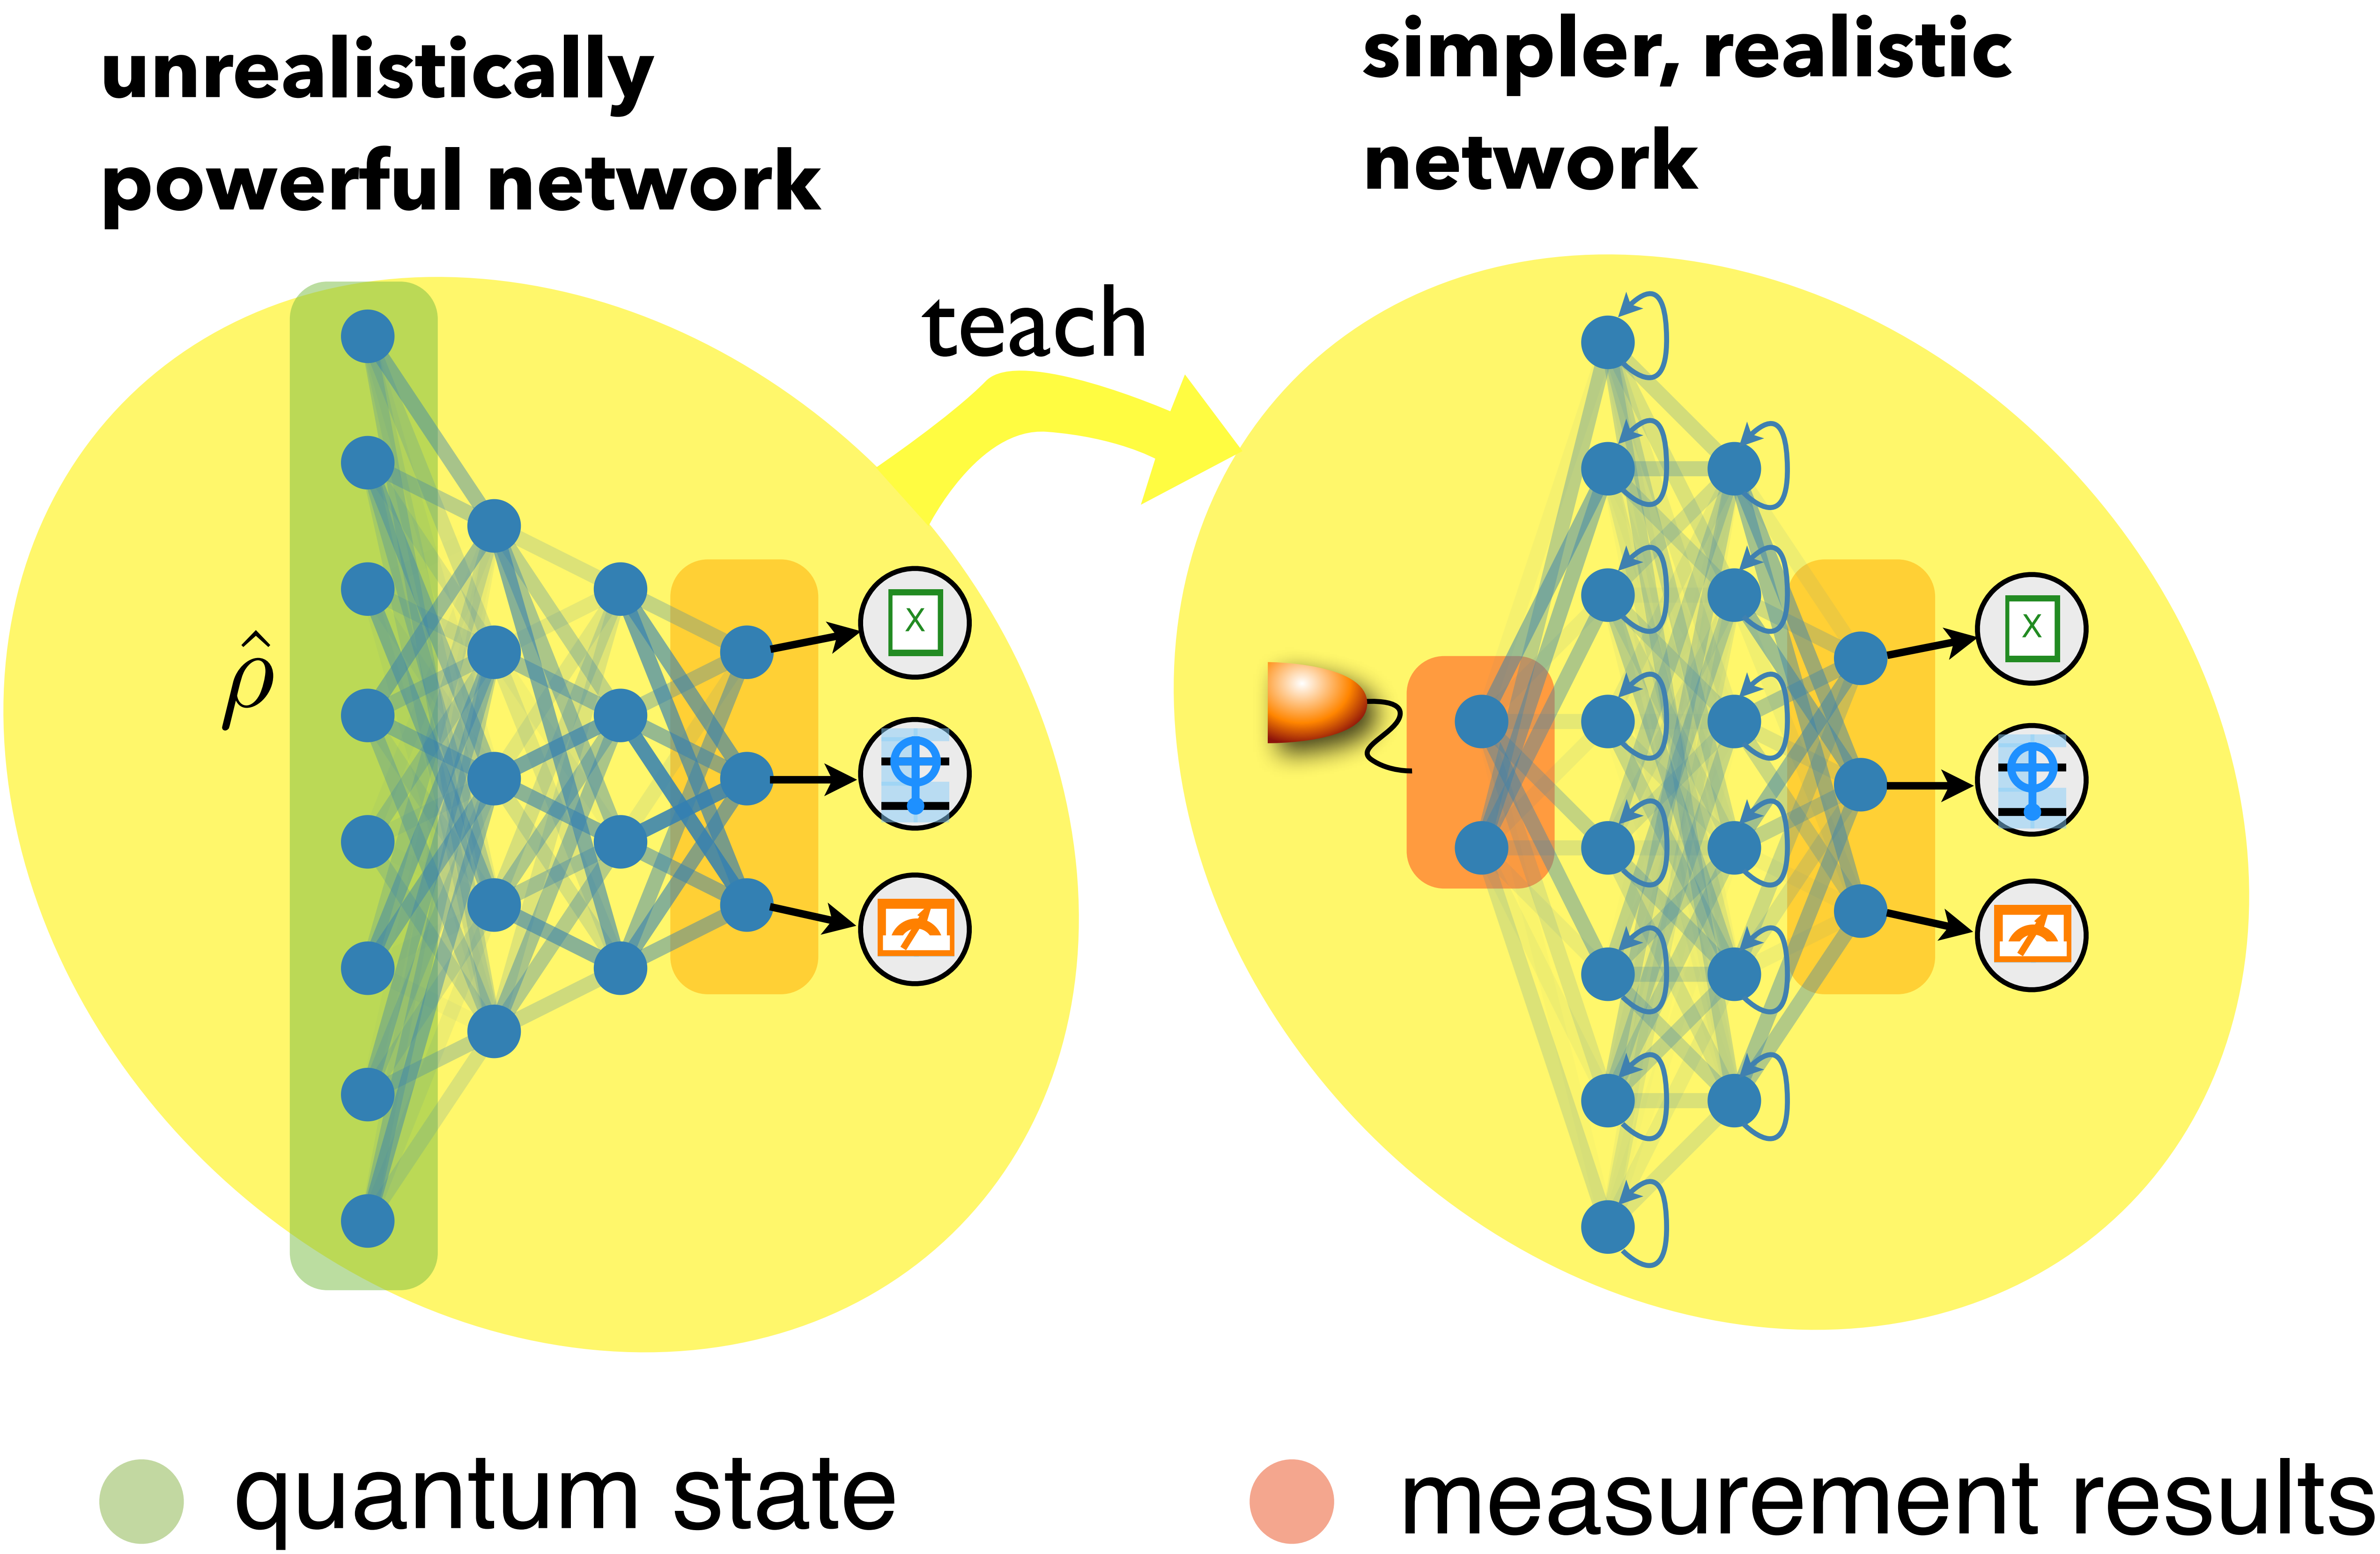

"Two-stage learning"

main advantage: flexibility
– applicable to many other physical settings

cavities

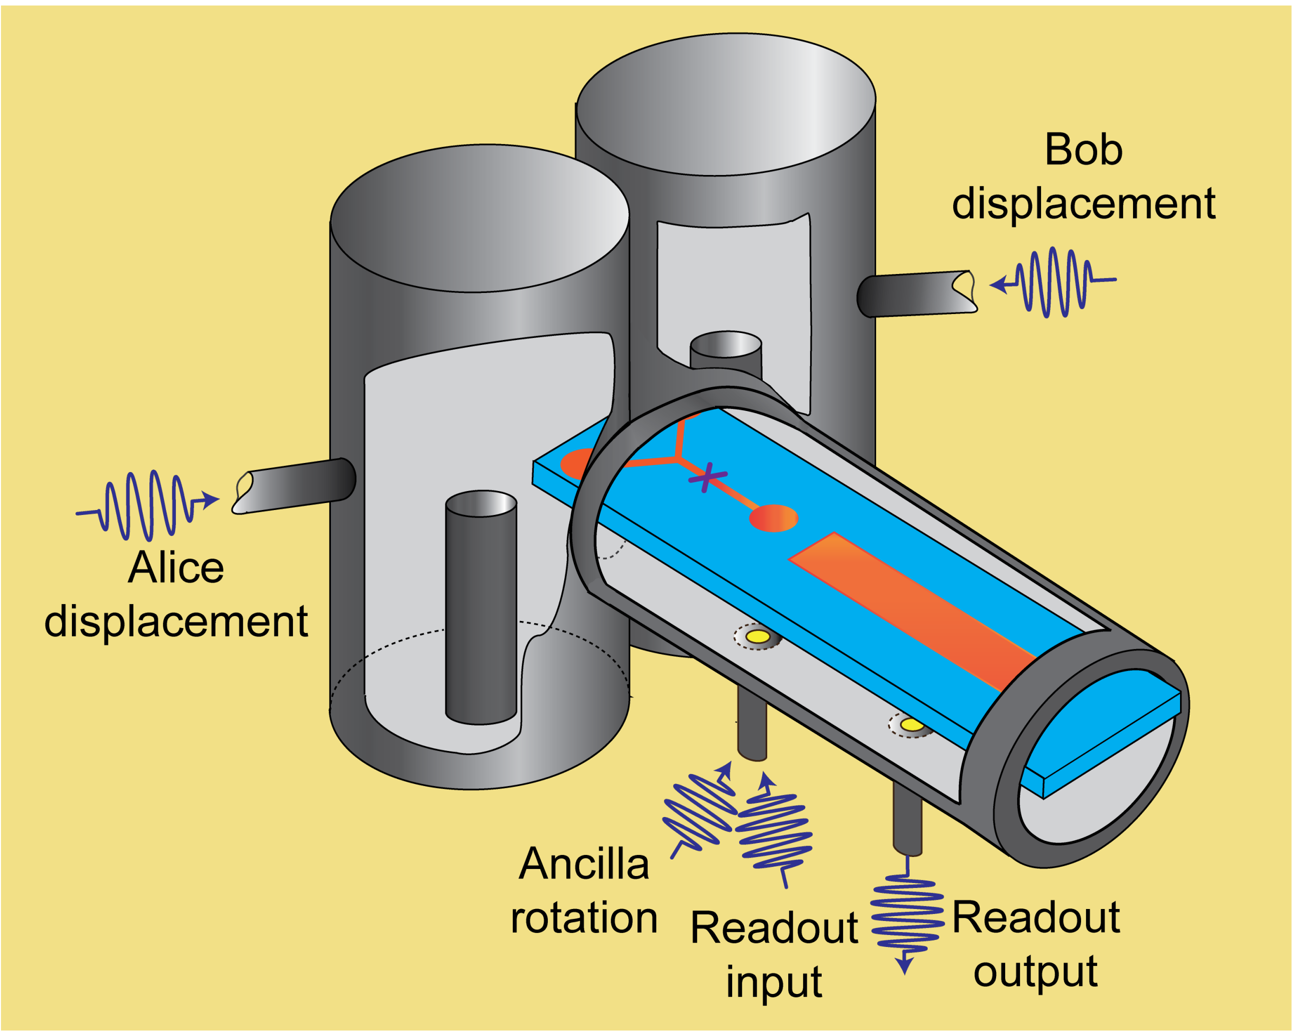

(Schoelkopf, Devoret lab 2016)

ion trap chips

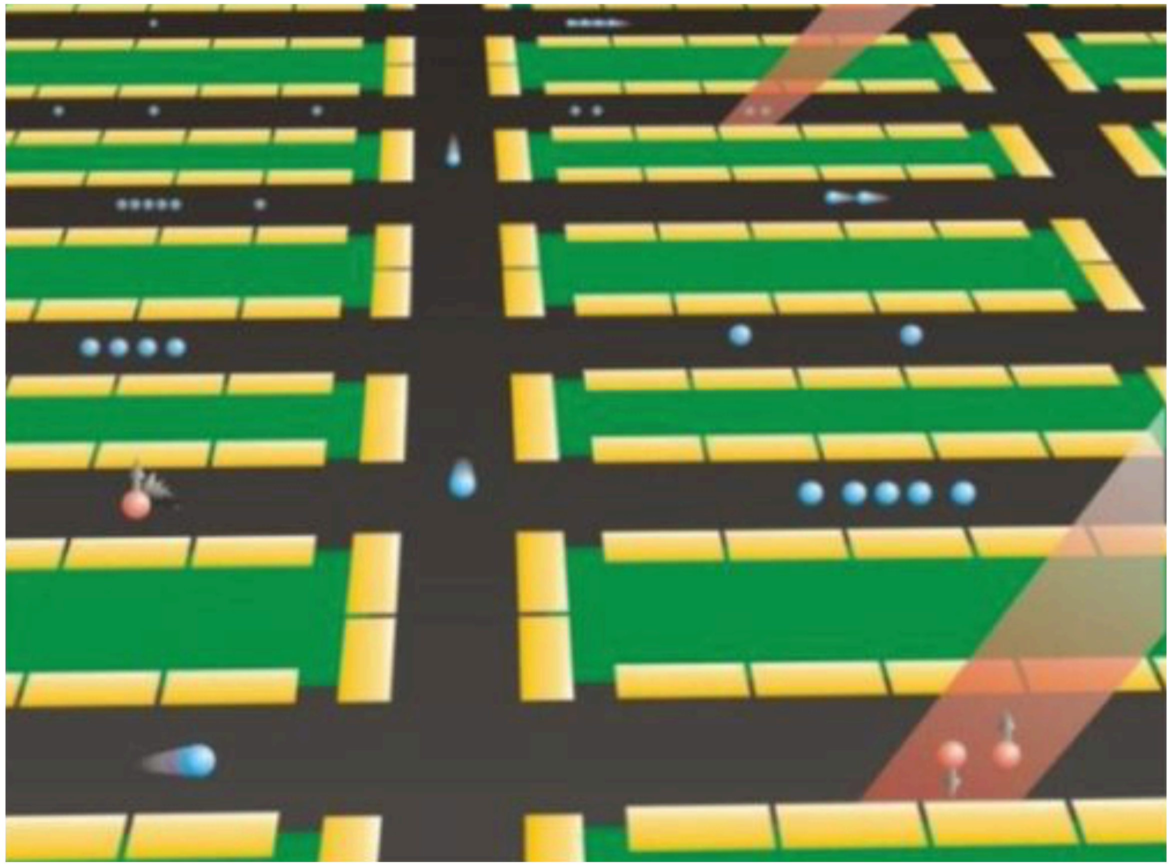

(Monroe, Kim Science 2013)

Currently: First steps towards
deep reinforcement learning on experimental platforms

example: optimizing control pulses for supercond. q. proc. (state for agent = quantum state tomography)

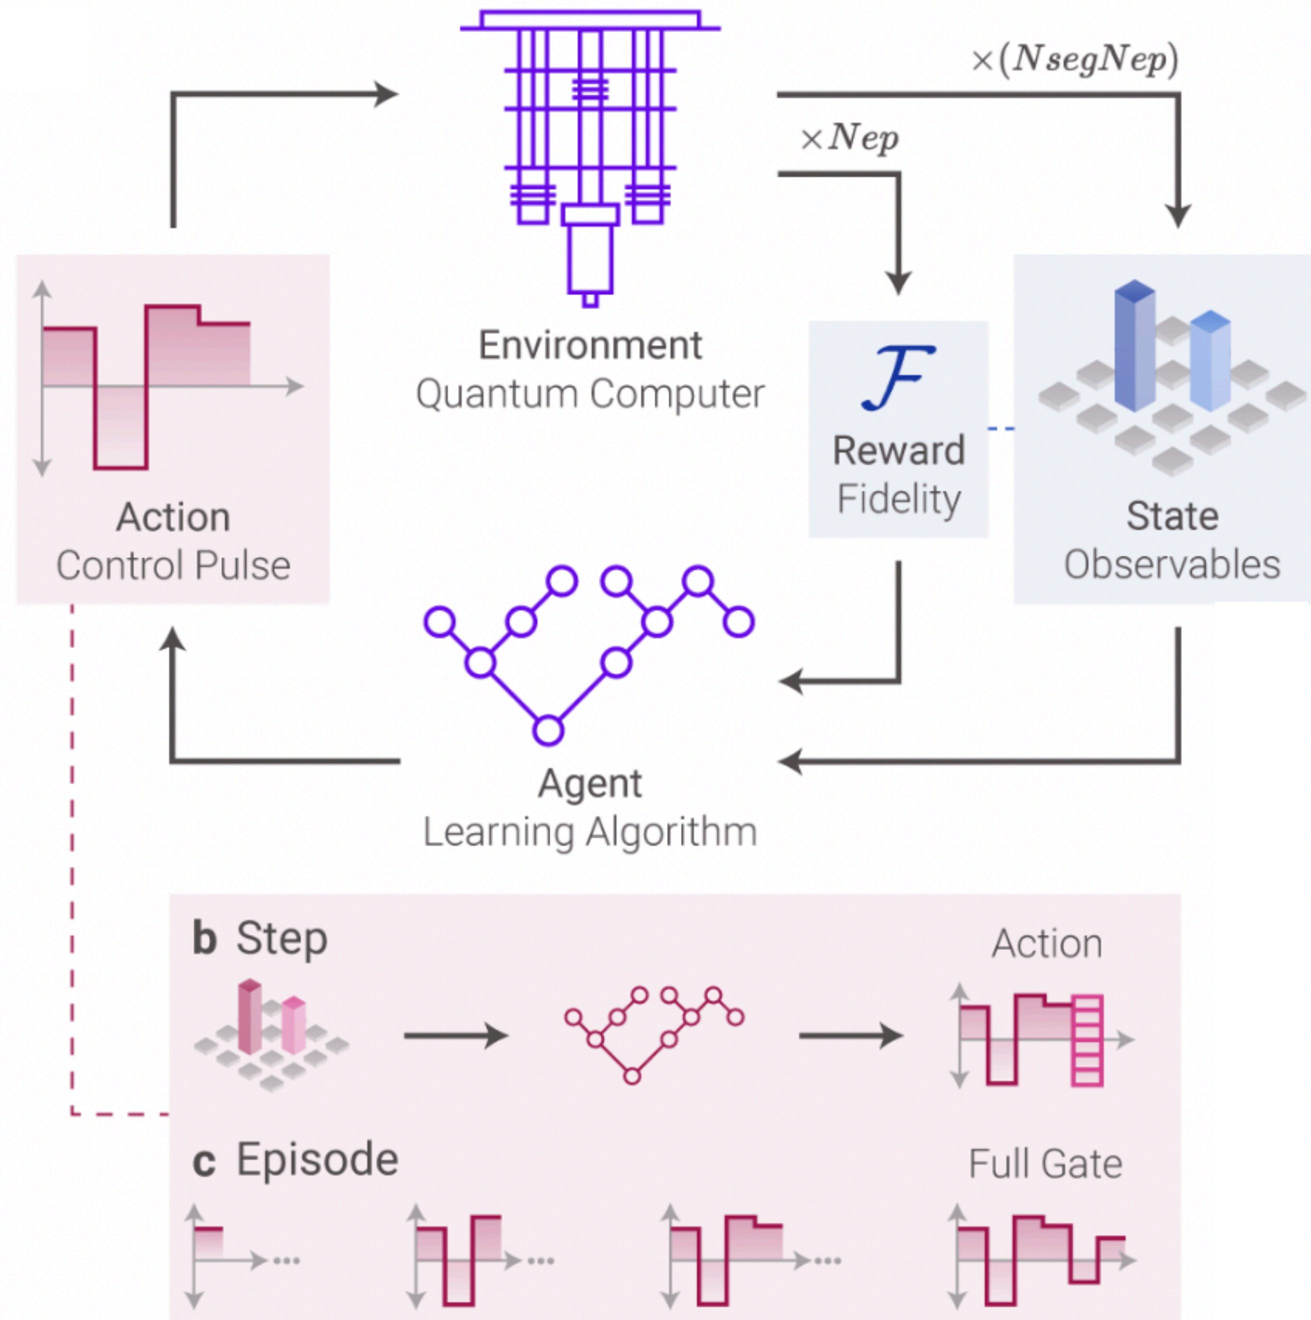

Baum et al. 2021

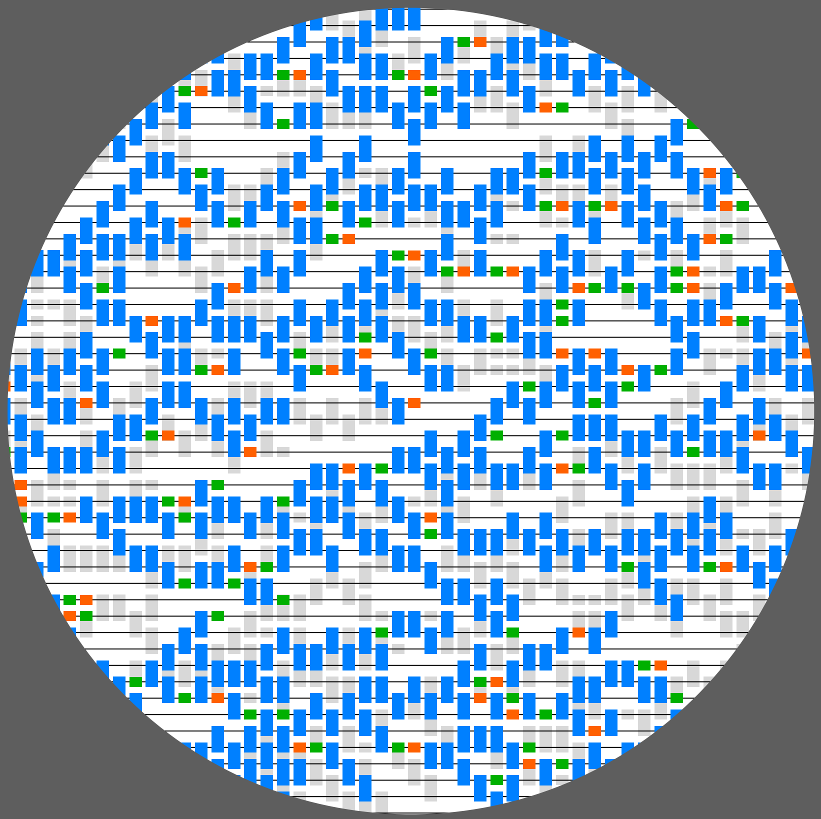

Quantum Circuit Optimization

work with Thomas Fösel, Murphy Yuezhen Niu, and Li Li (Google Research)

arXiv 2103.07585

Quantum Circuit Optimization: reduce gate count / depth / etc. !

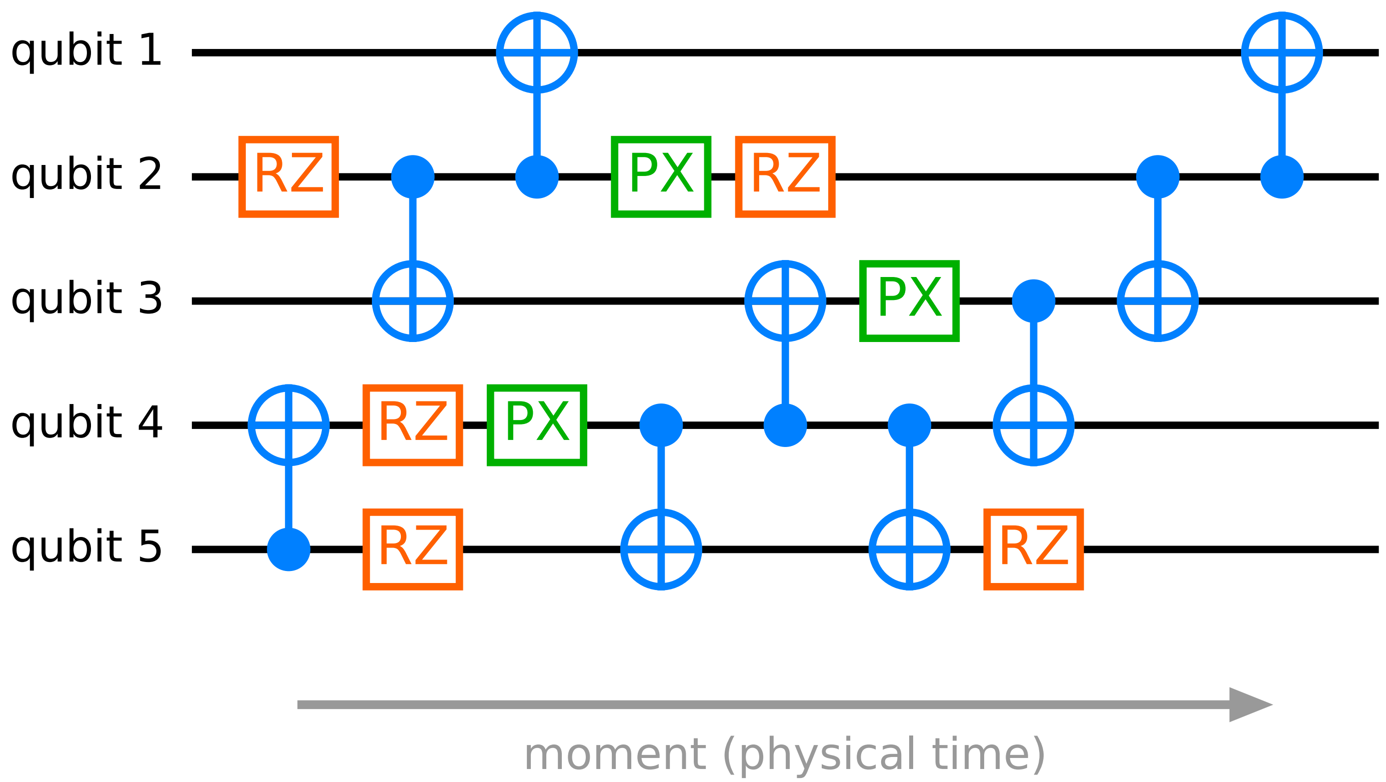

NISQ devices: Quantum Circuit Optimization critical …and needs to be hardware-dependent [not on an abstract level designed for large-scale fault-tolerant circuits]

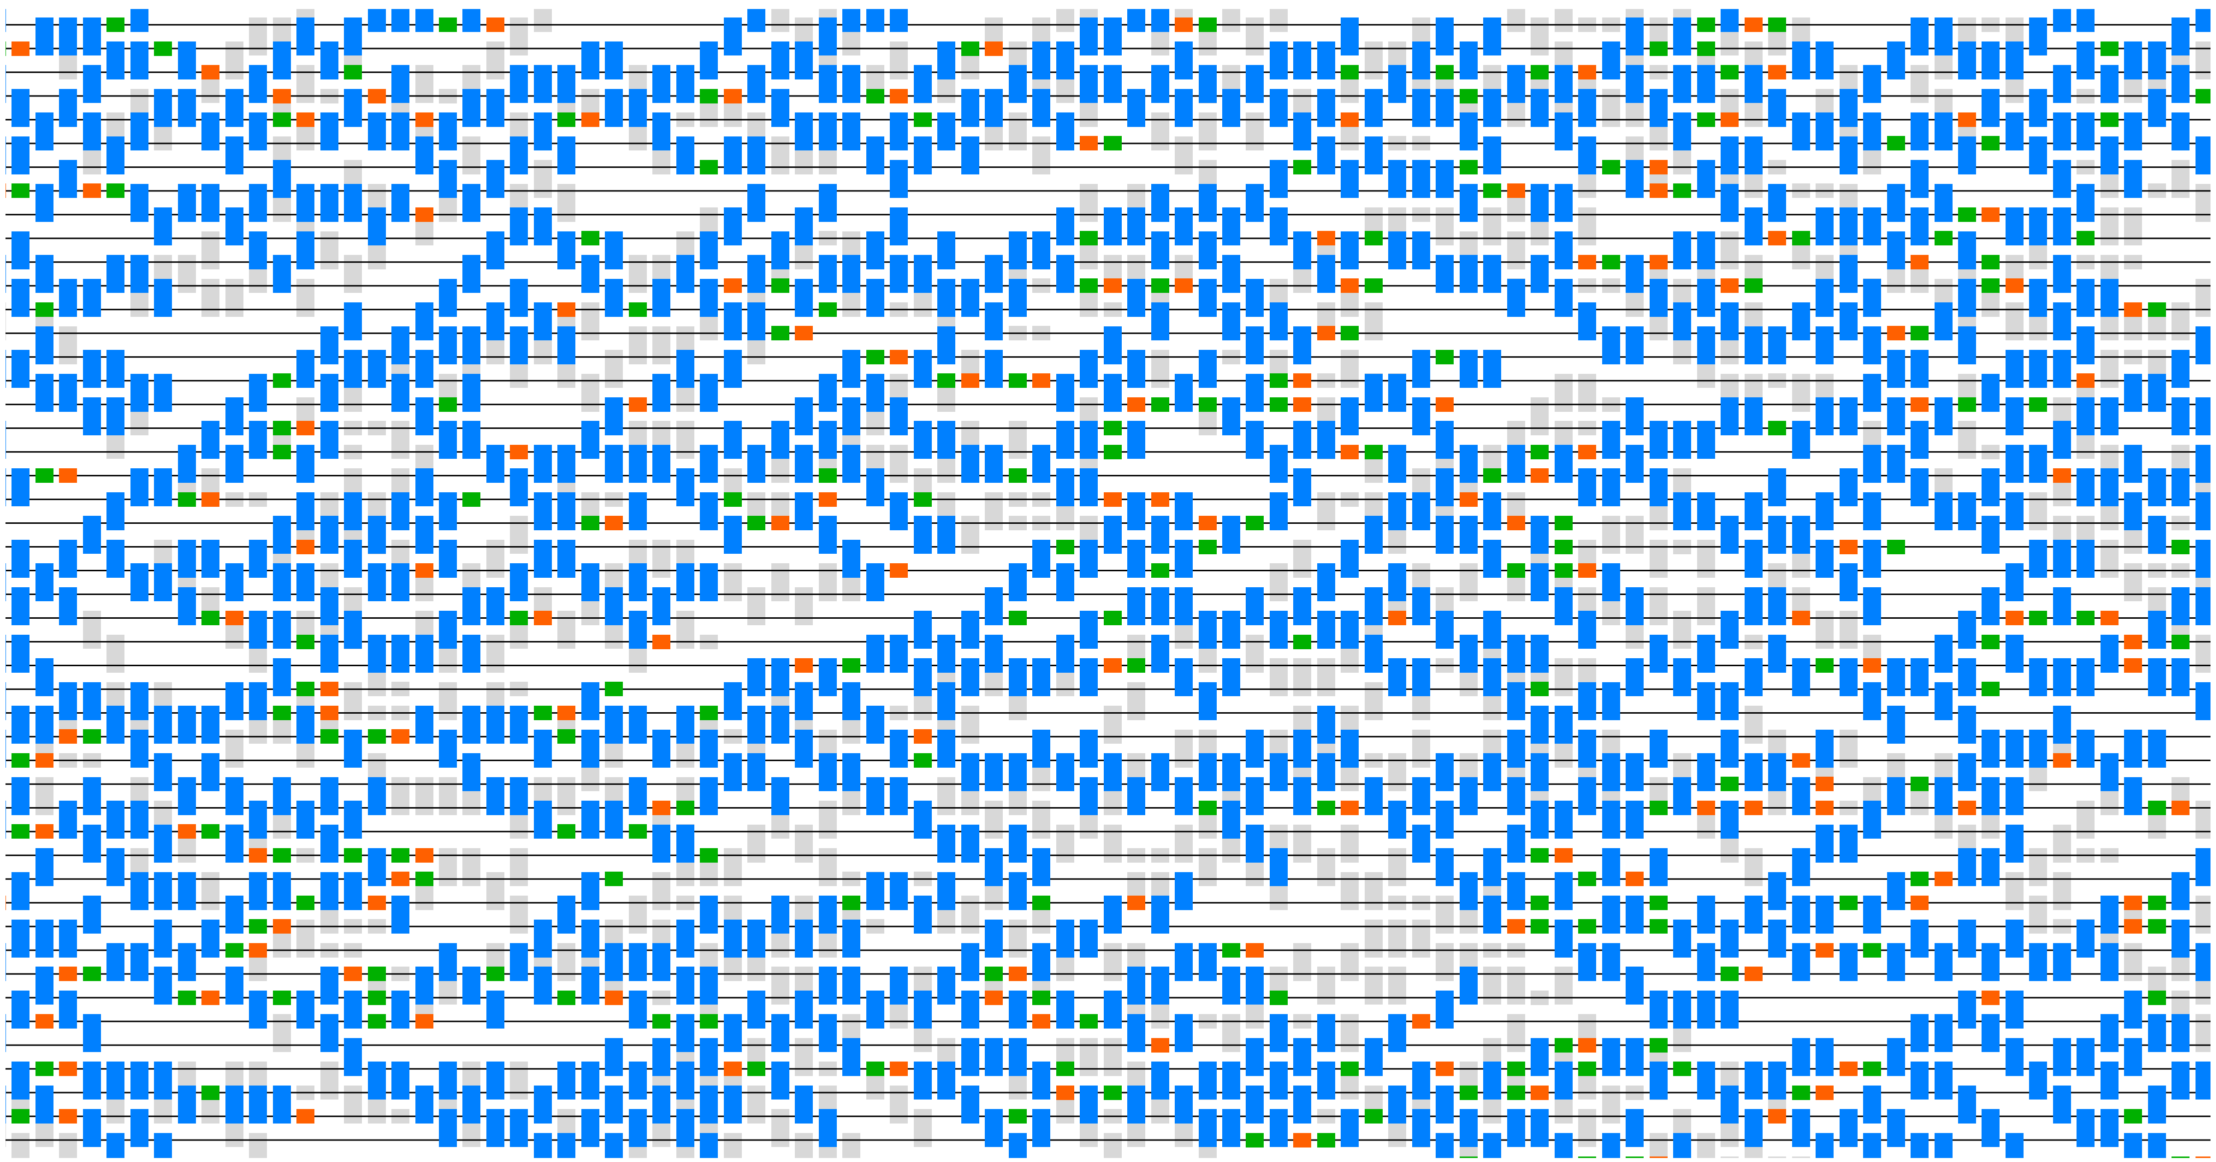

Transformation rules

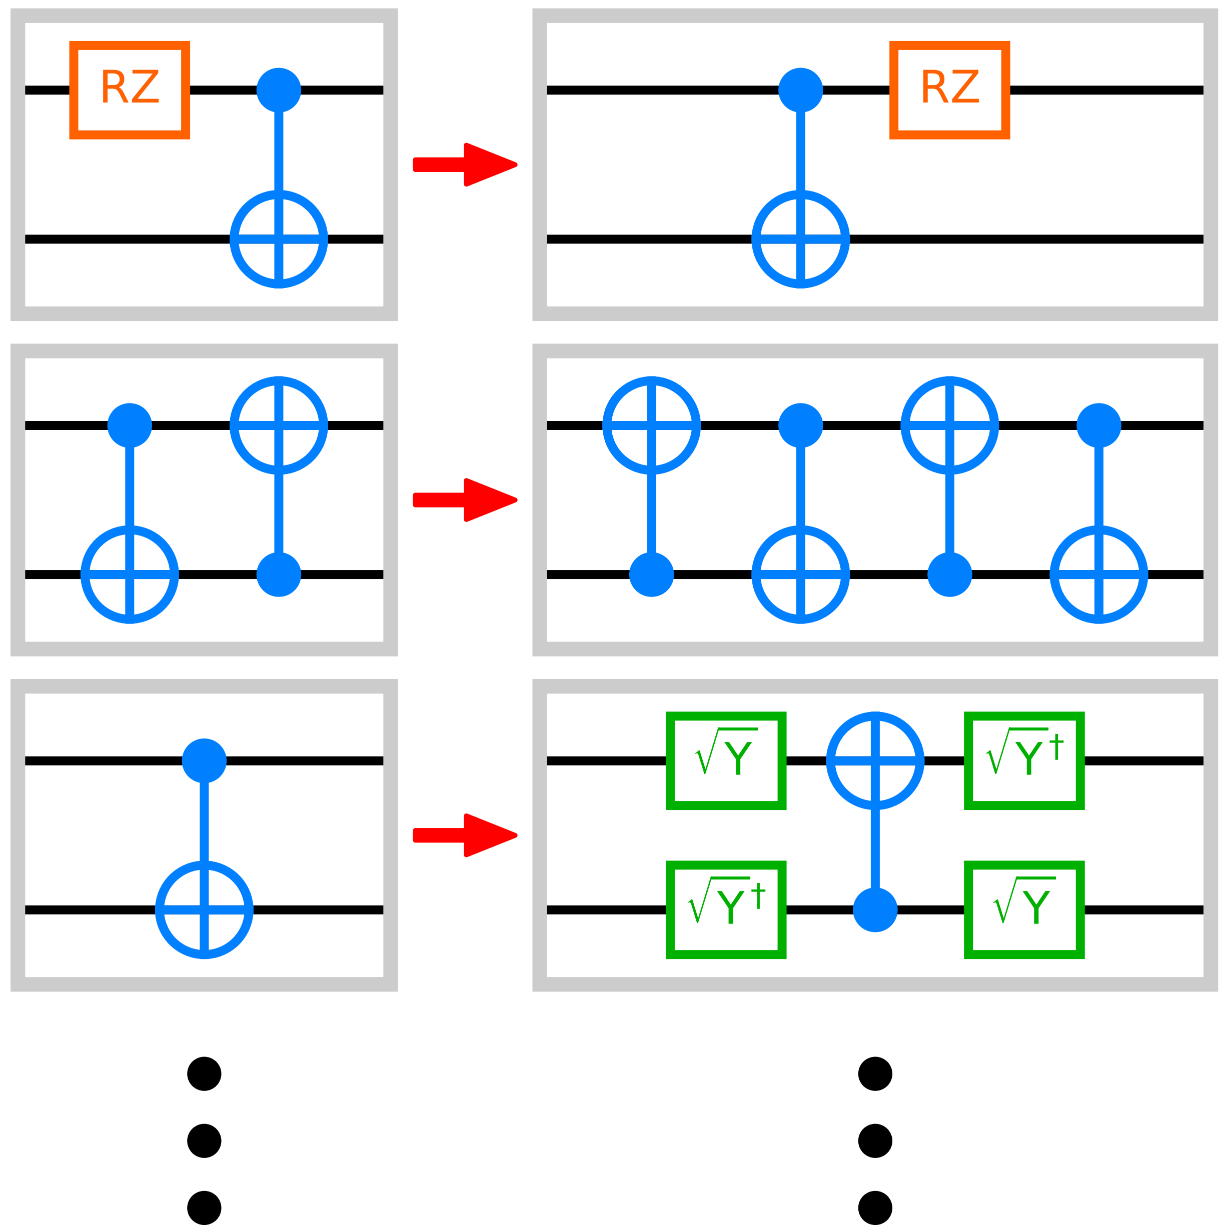

Deep reinforcement learning approach

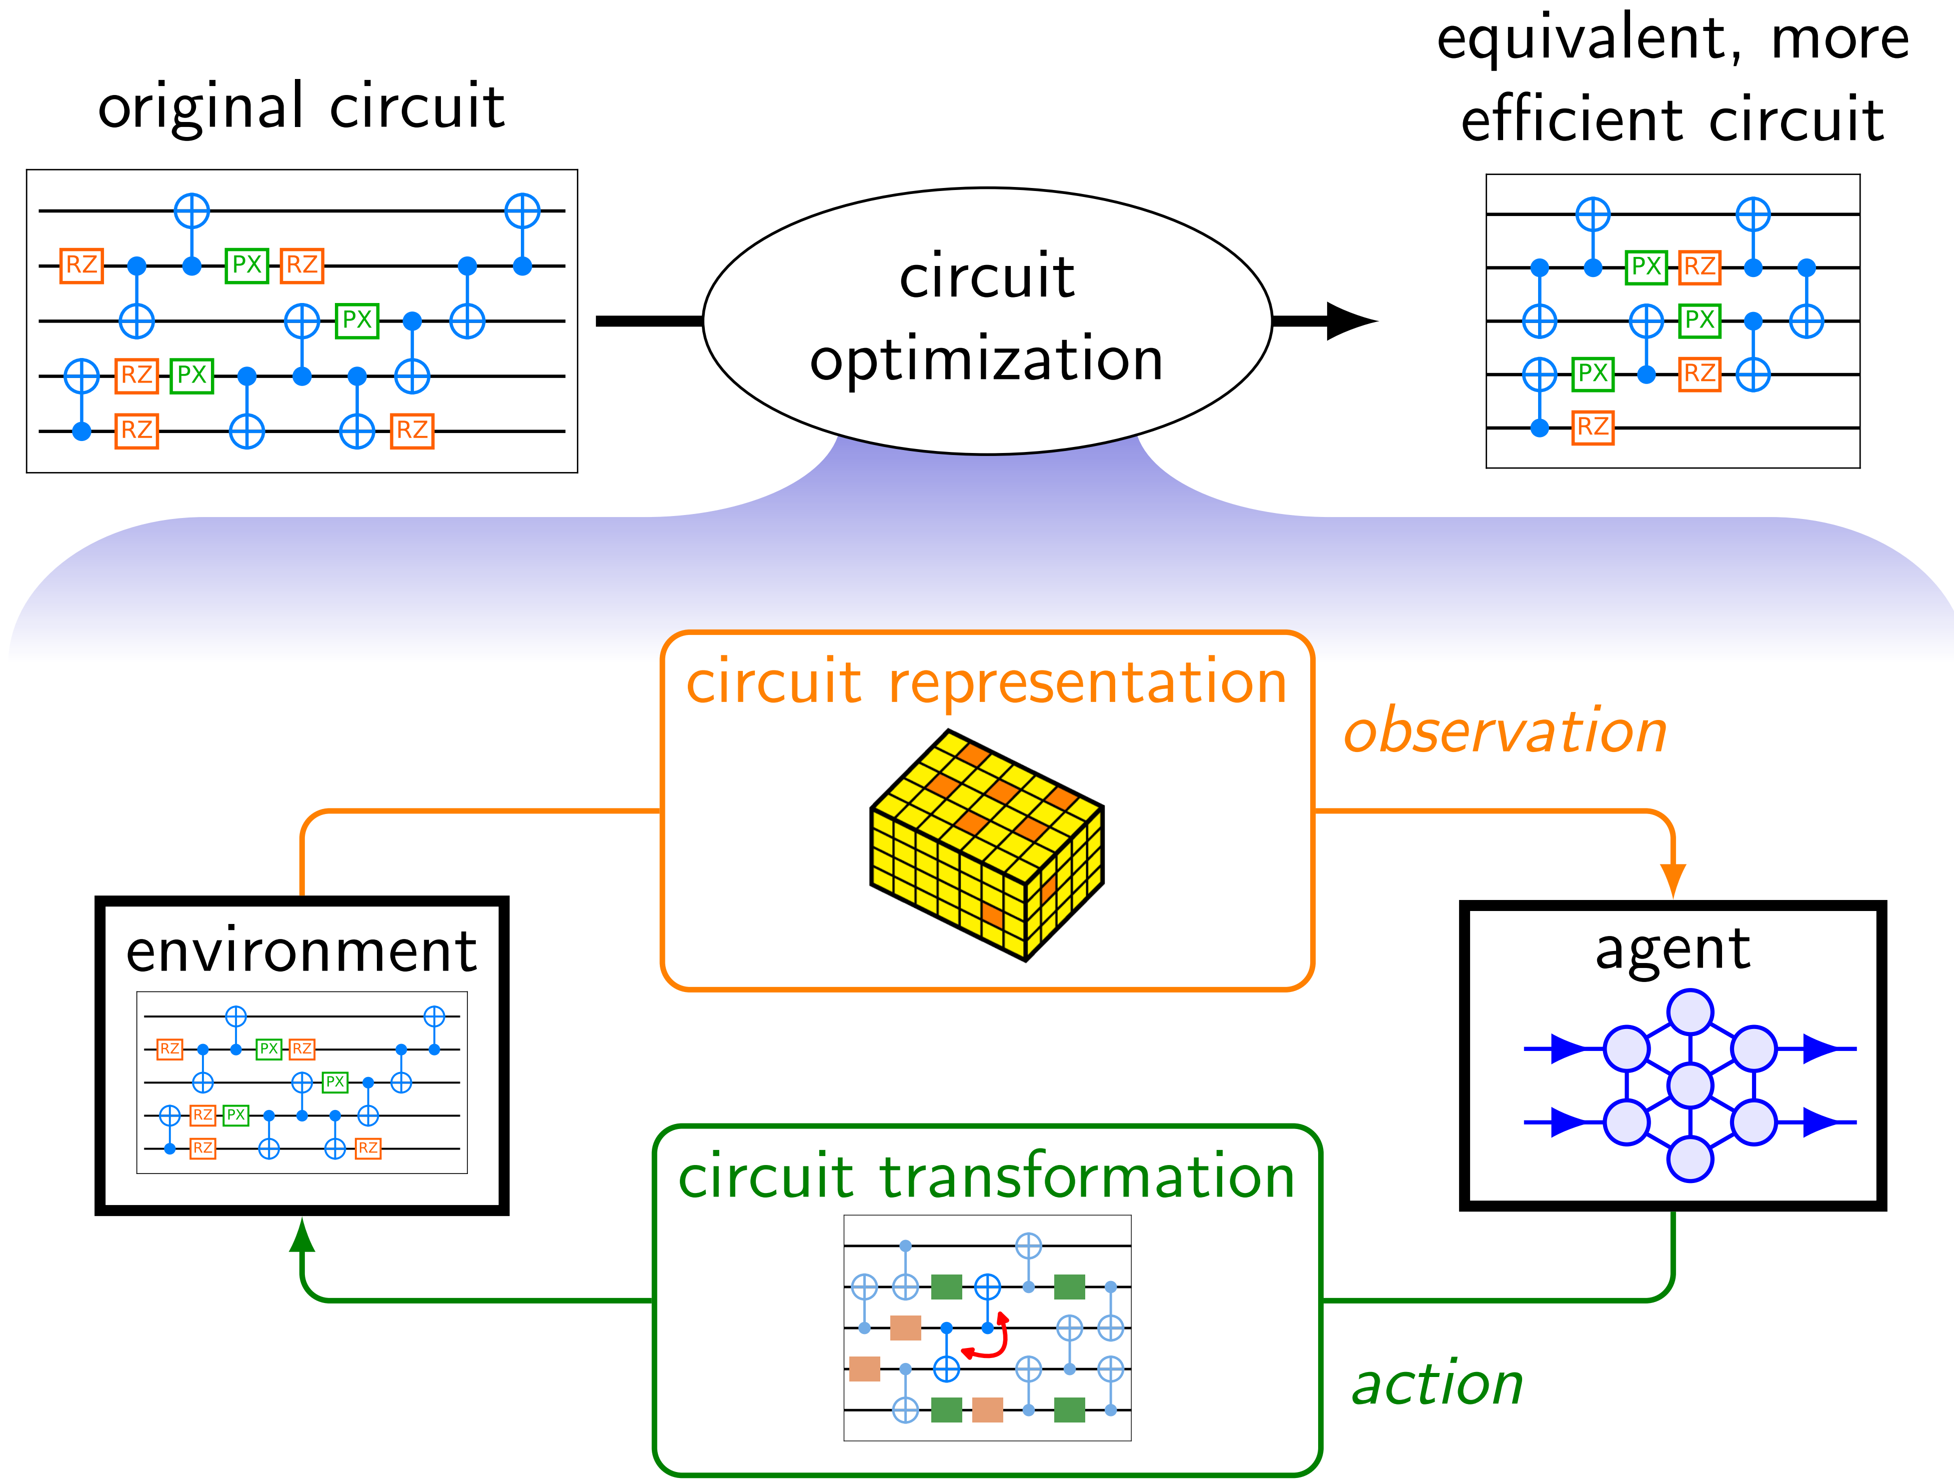

hardware-efficient, cross-platform, autonomous, reliable

Choices: States, Agent, Actions, Rewards

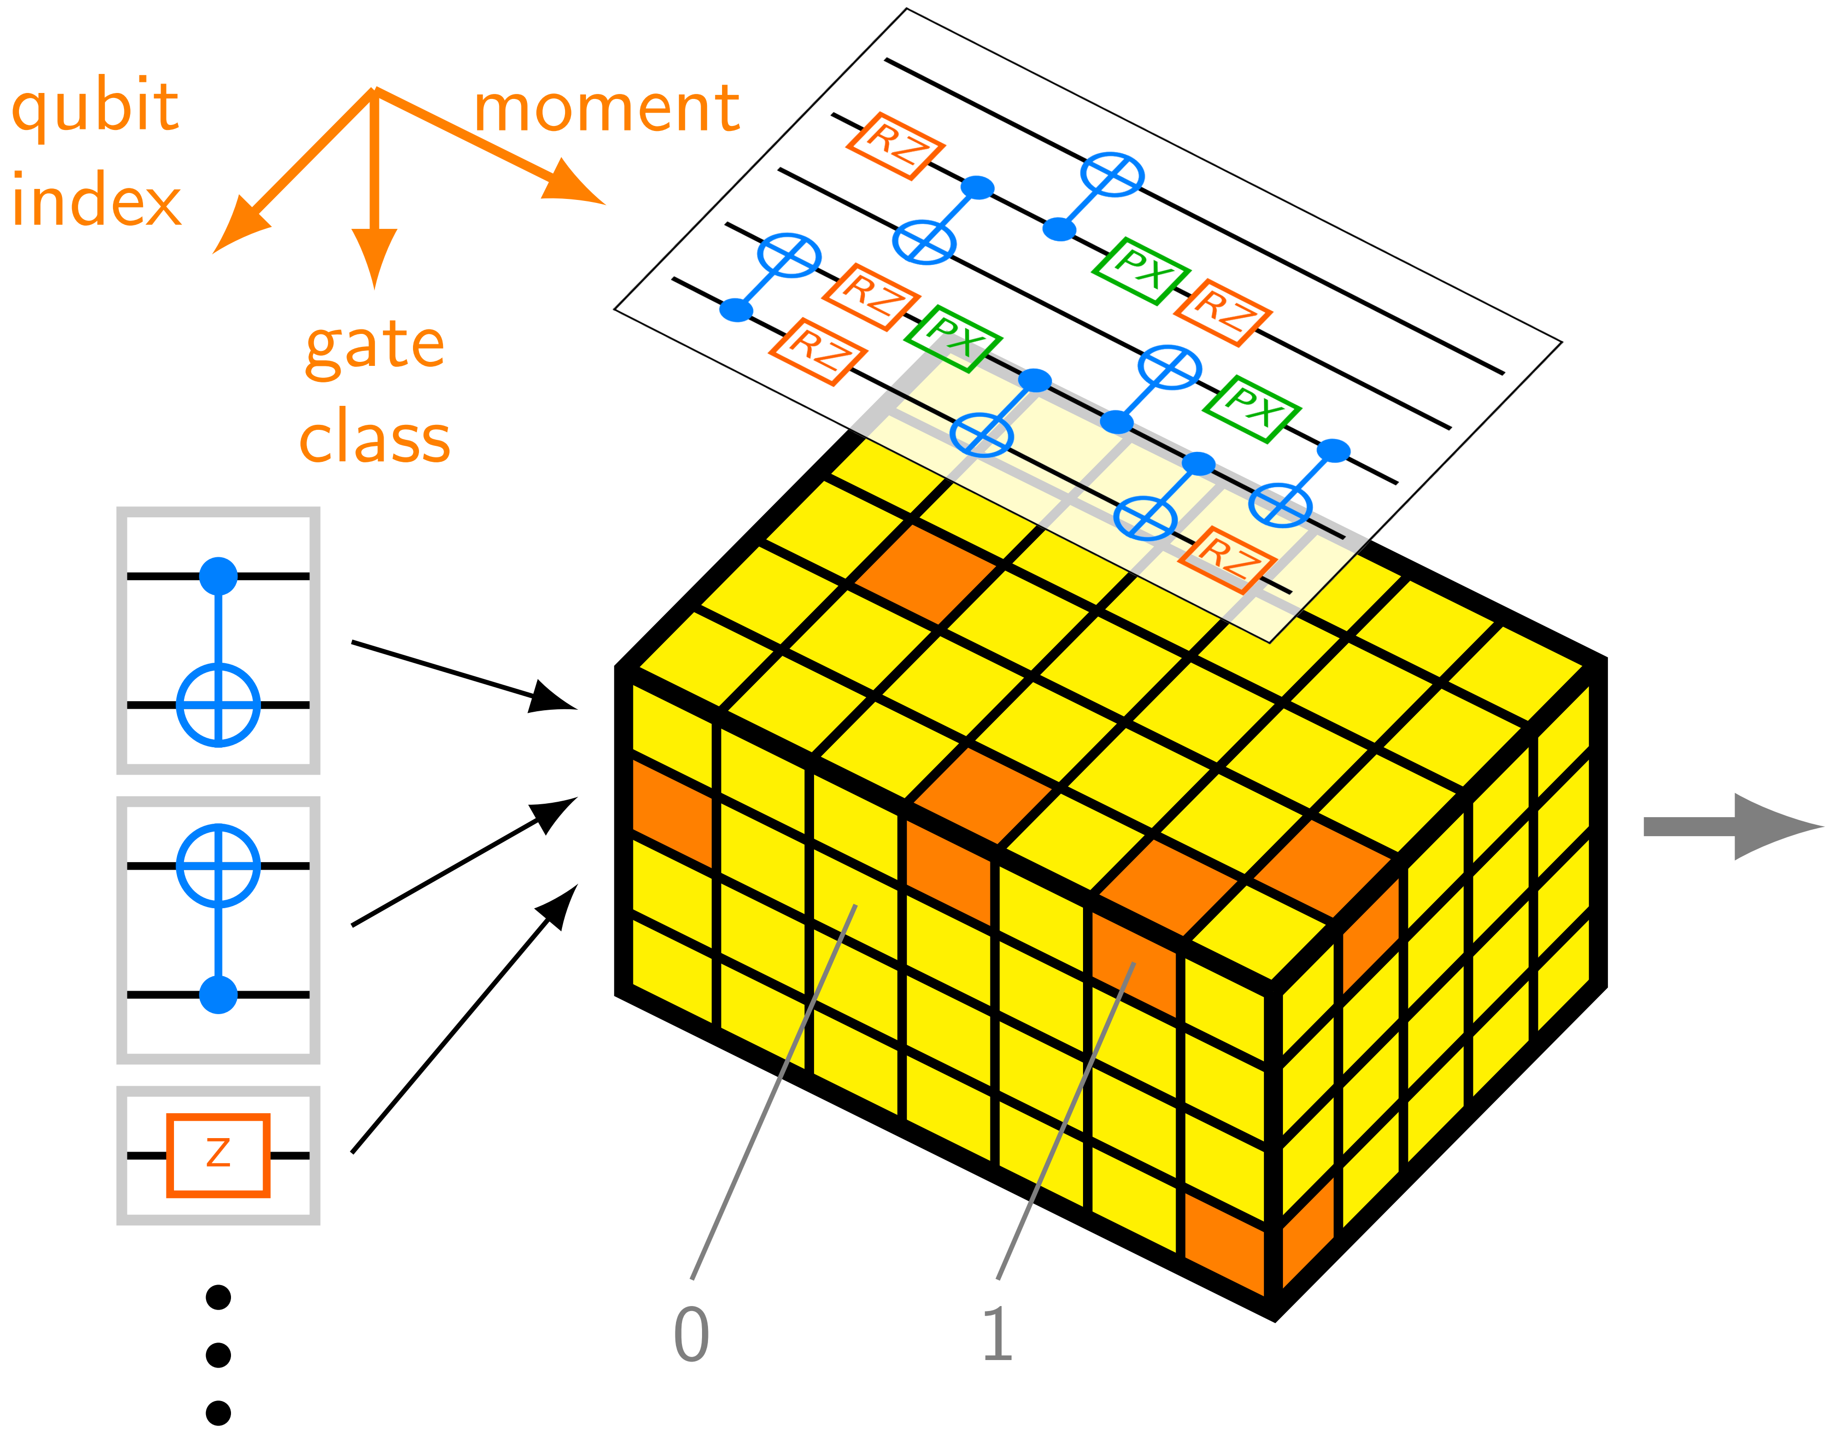

observation for circuit=state s

Choices: States, Agent, Actions, Rewards

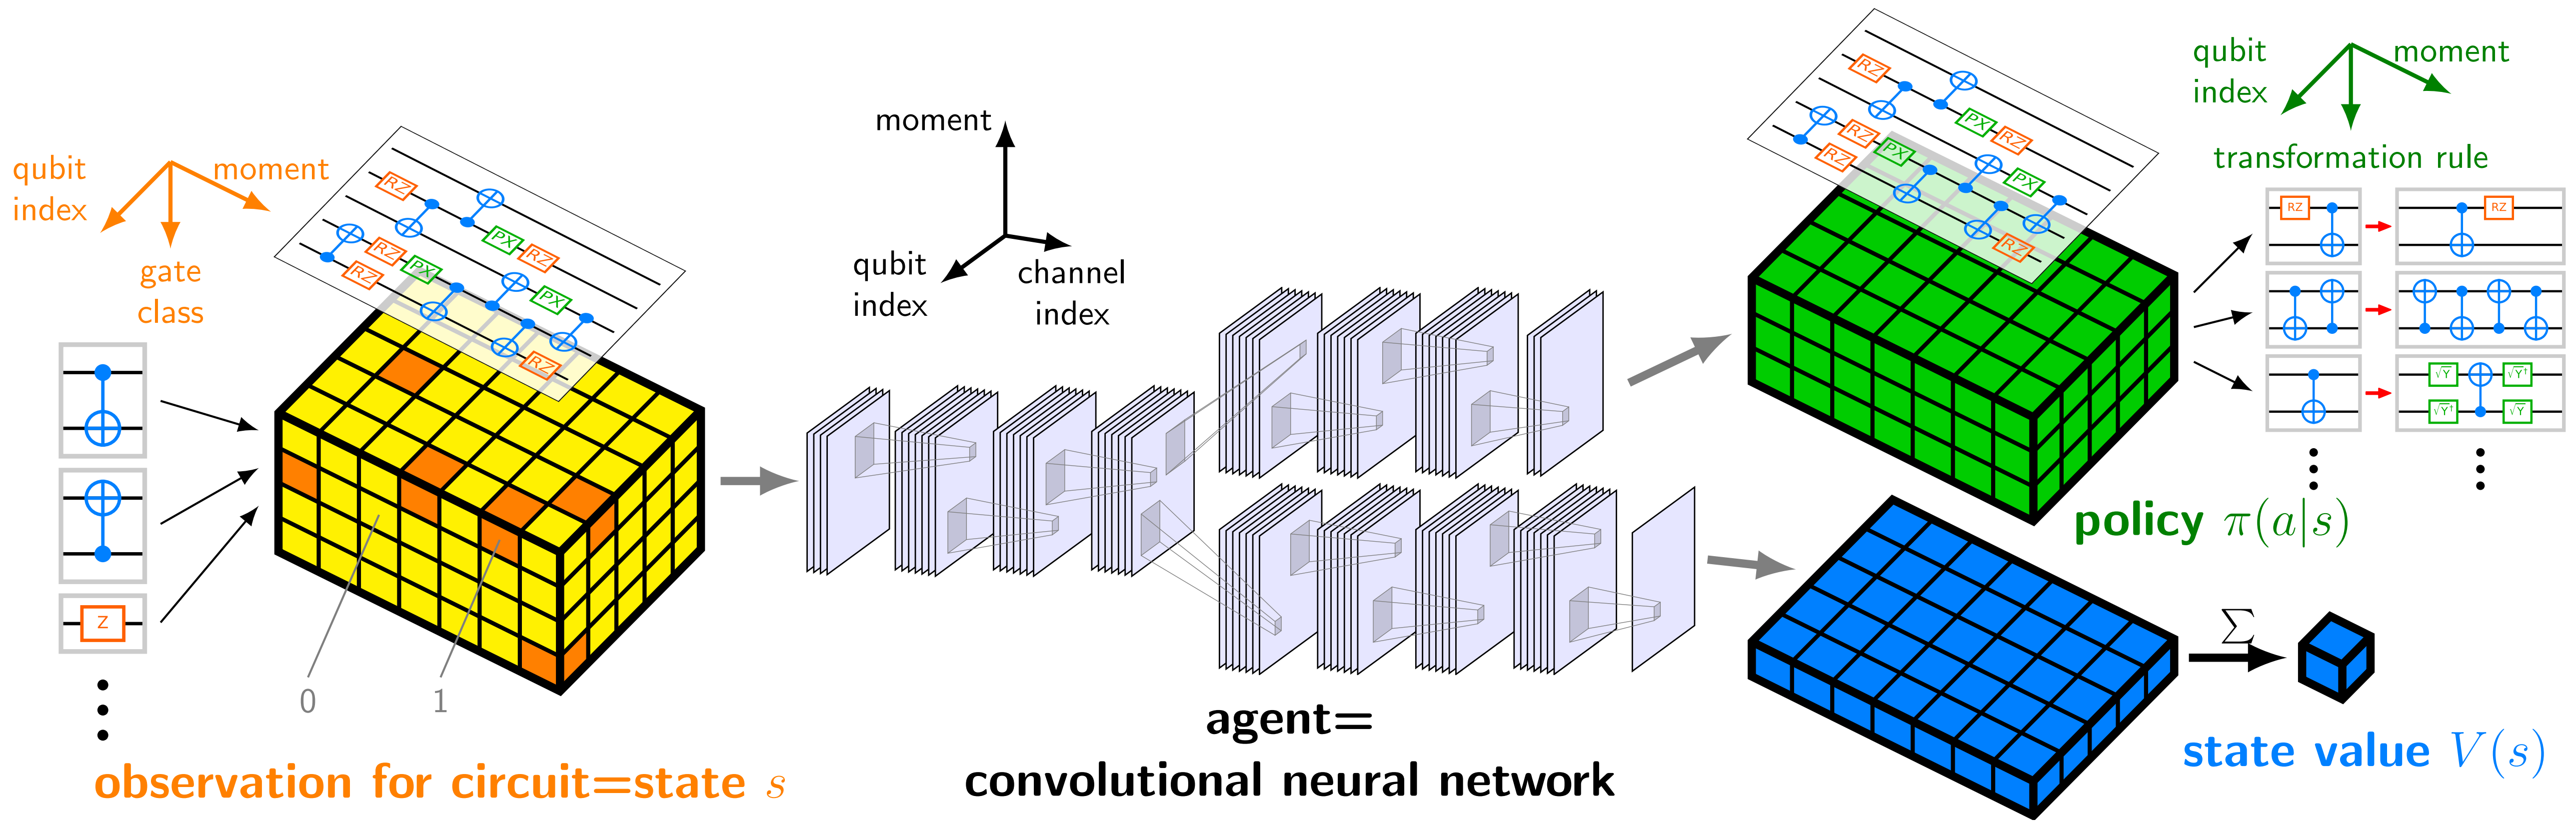

Technique: Advantage Actor Critic (namely: PPO)

Choices: States, Agent, Actions, Rewards

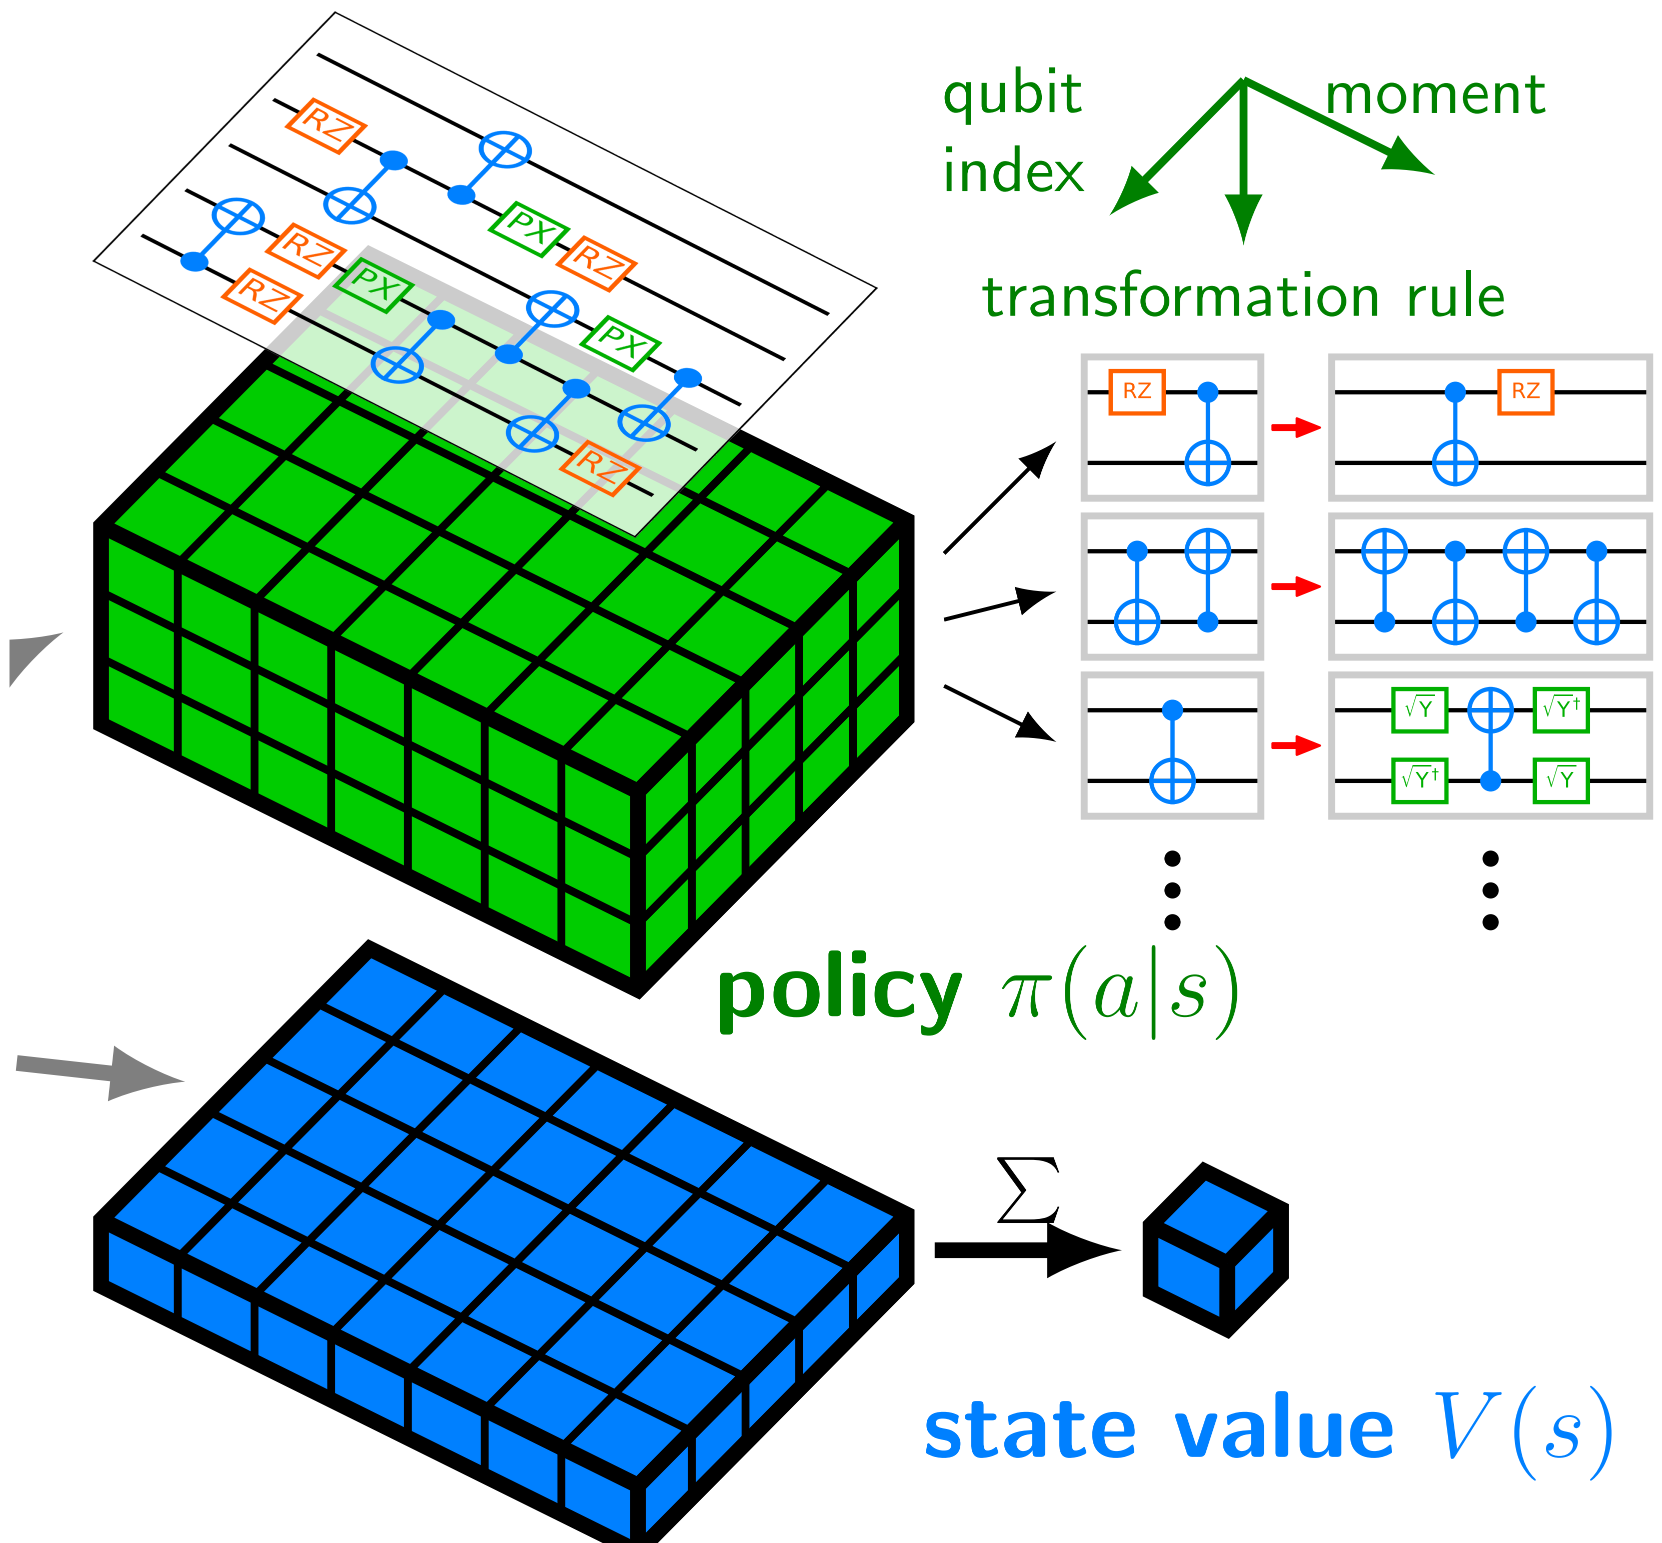

Choices: States, Agent, Actions, Rewards

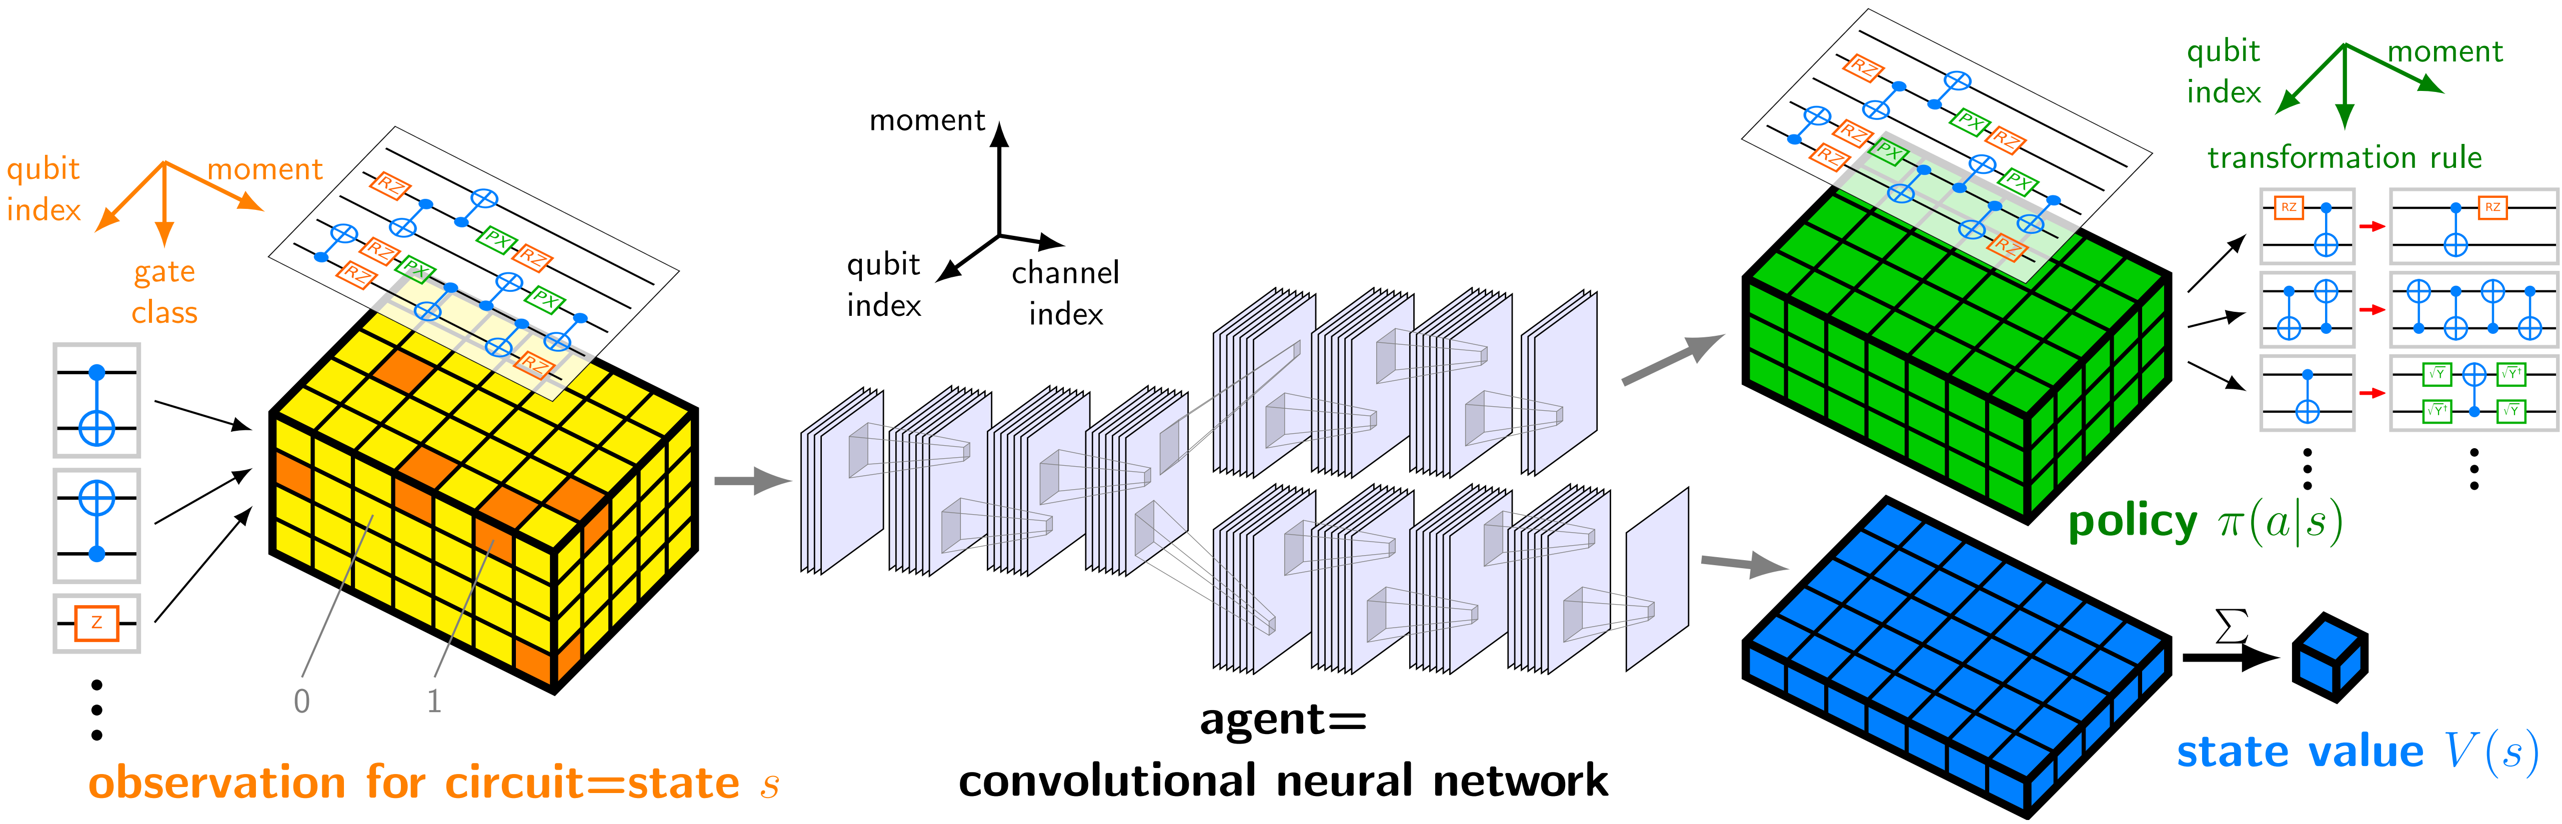

Reward: reduction in gate count, depth, or combination (possibly: gate-dependent, decoherence estimate, …)

Training on Random Circuits

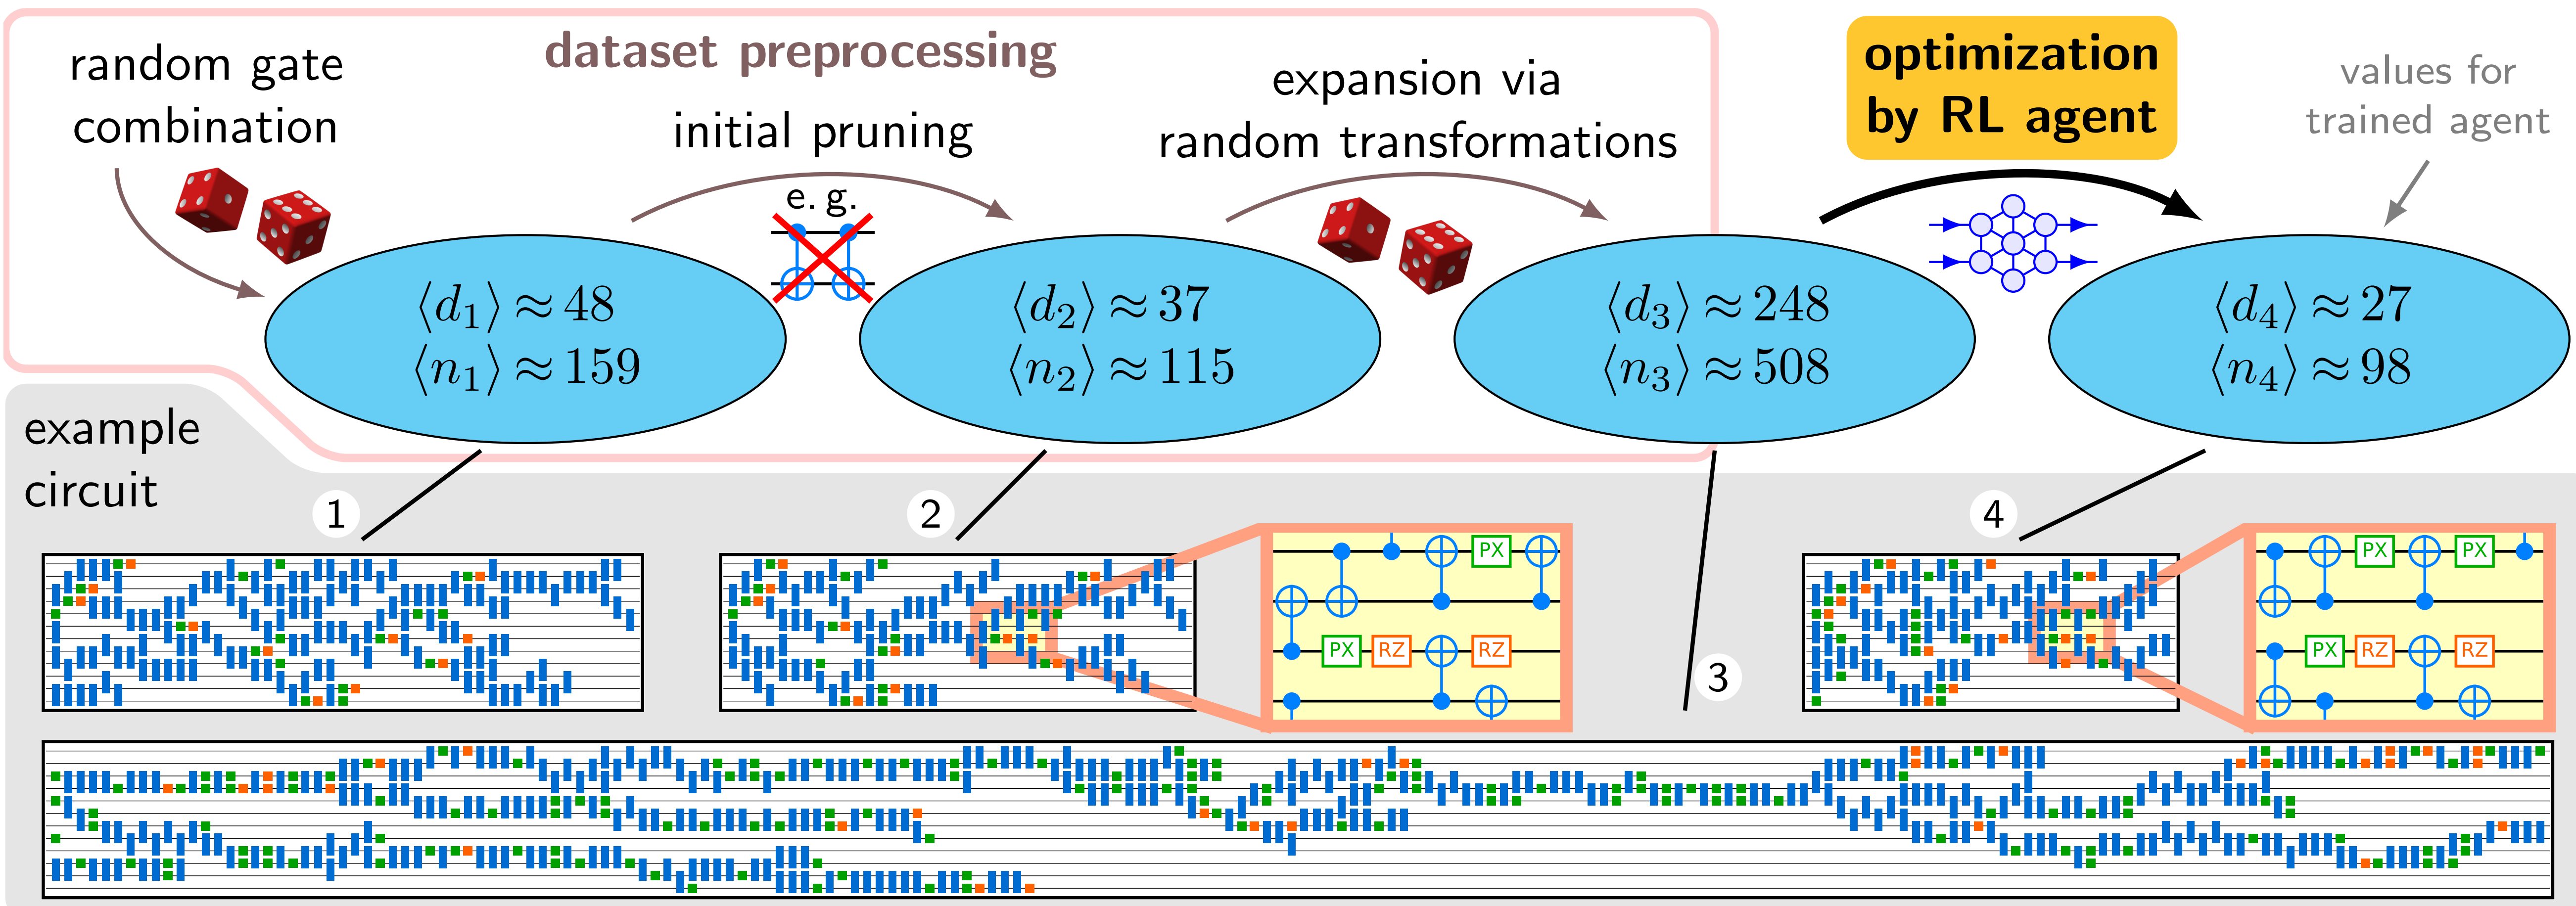

Training on Random Circuits: Progress

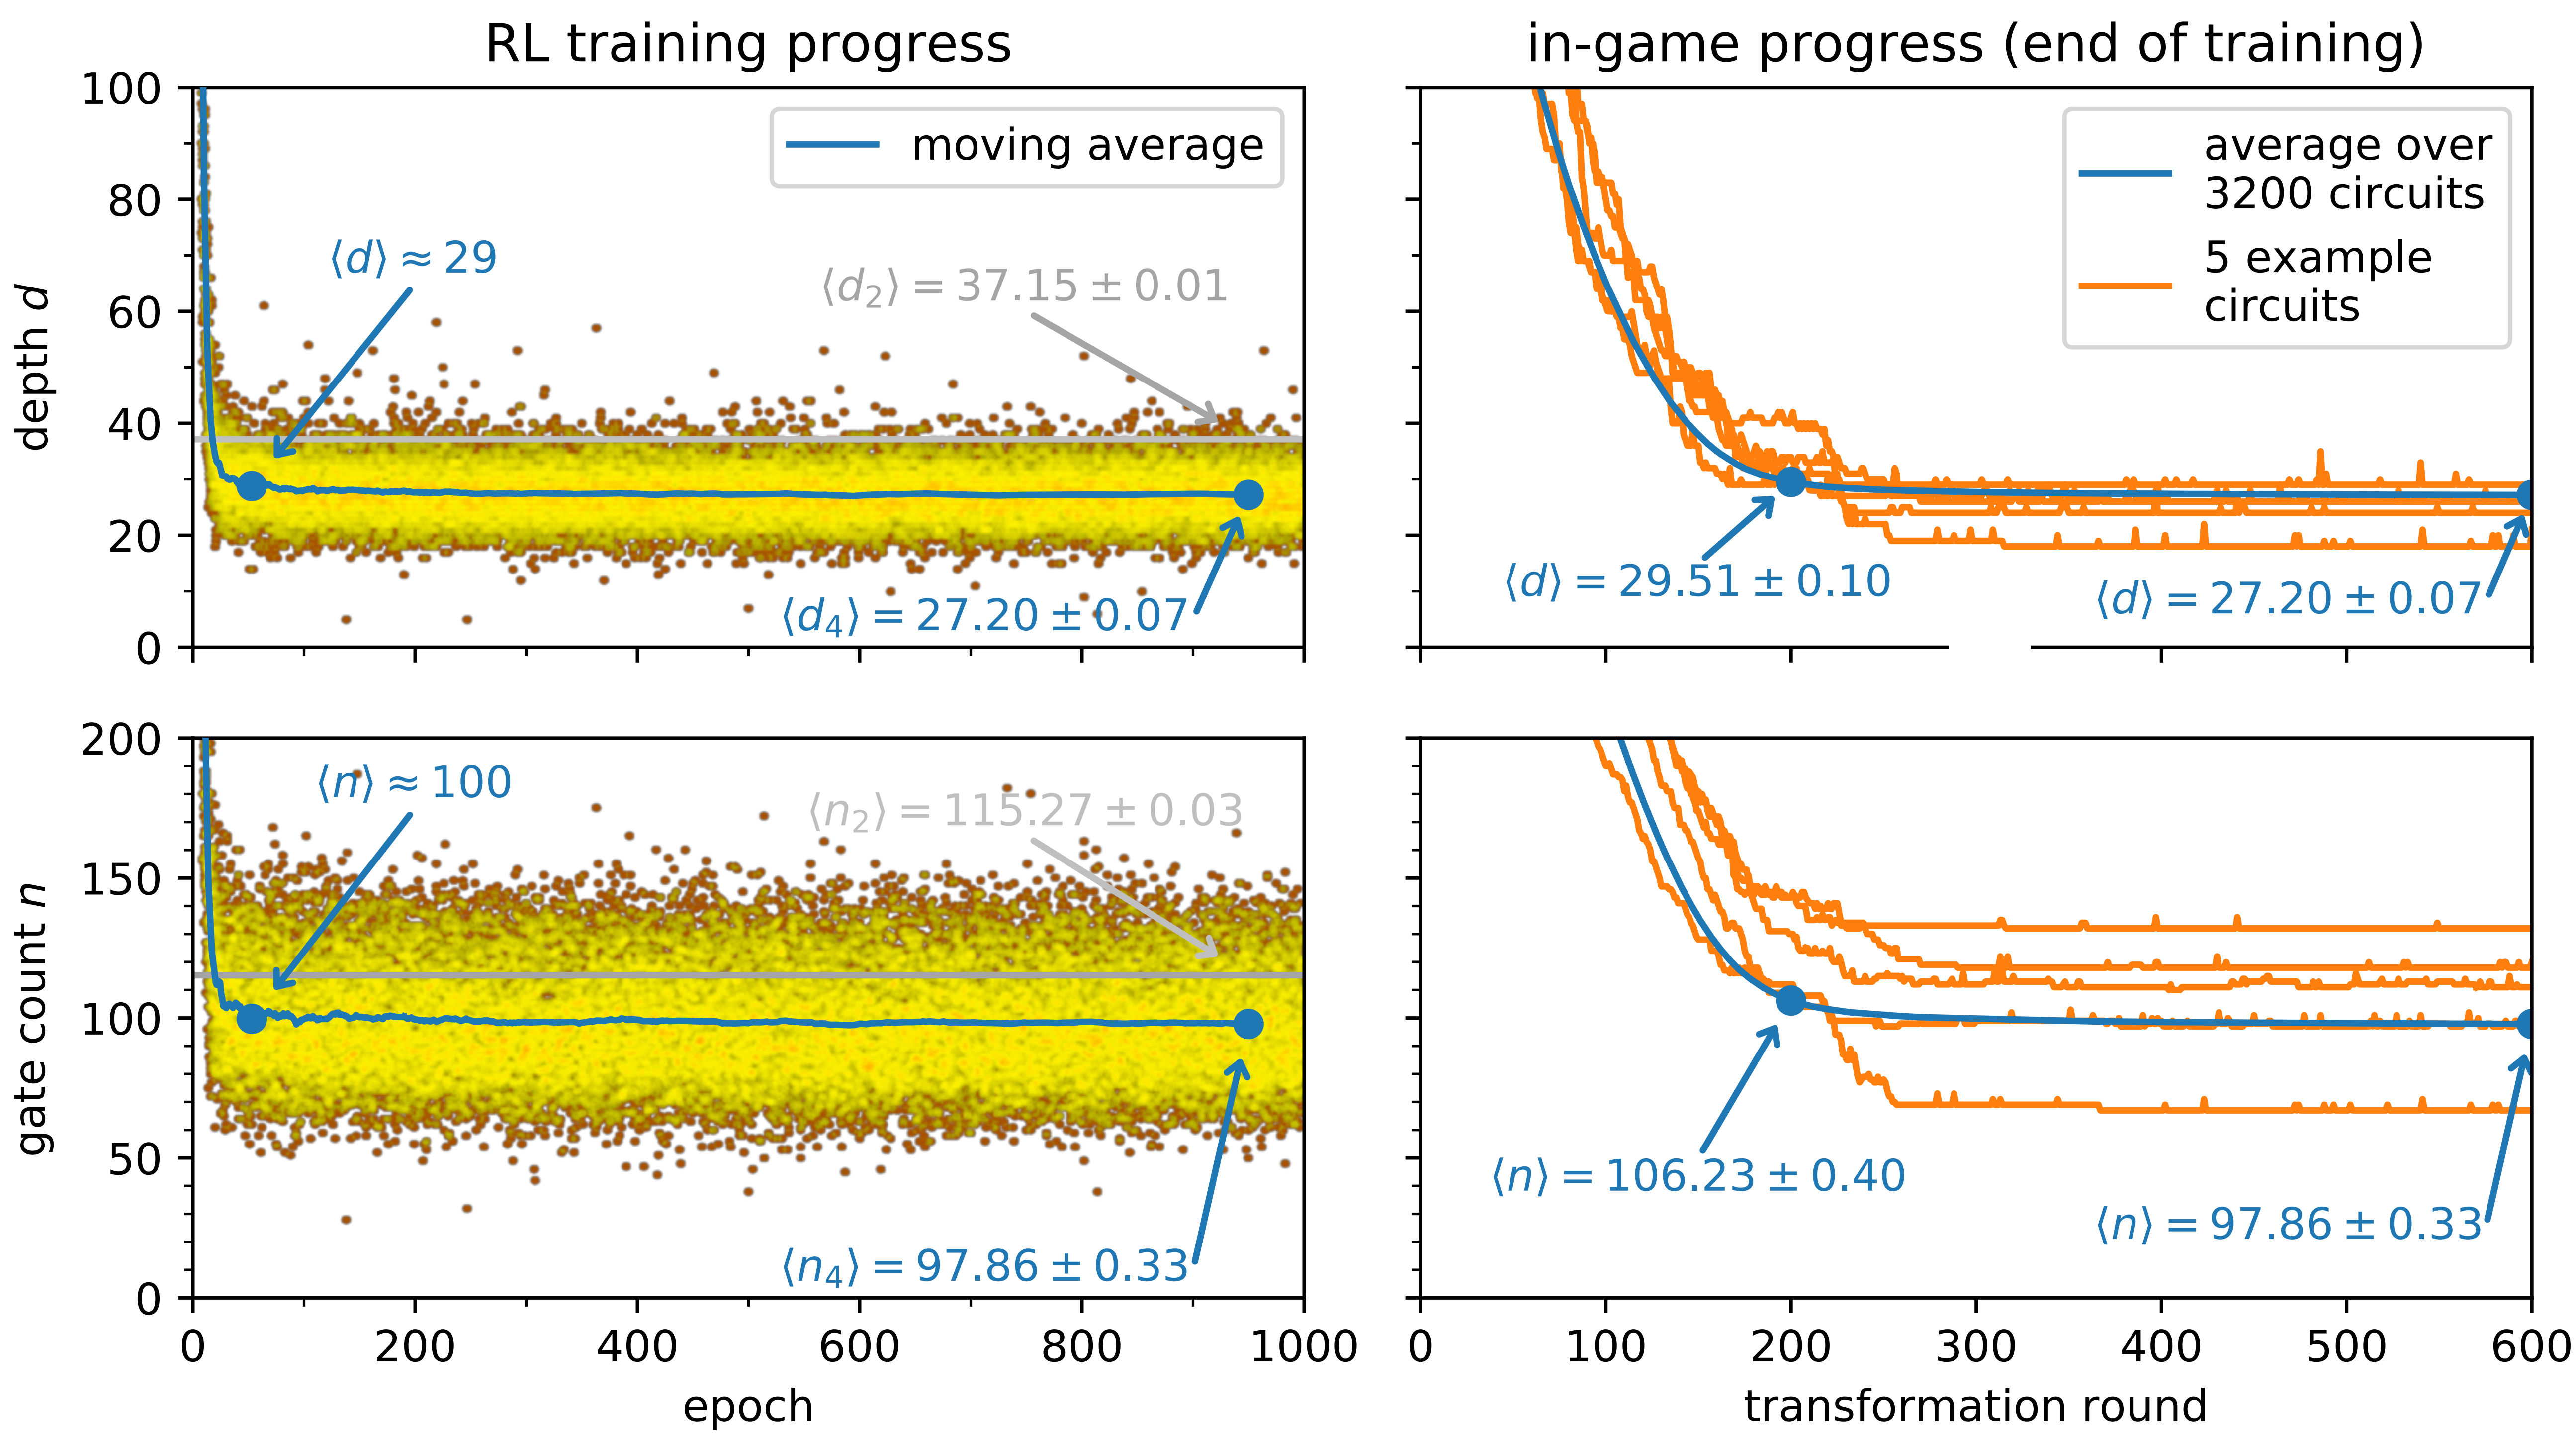

1 epoch = 32 episodes

Performance

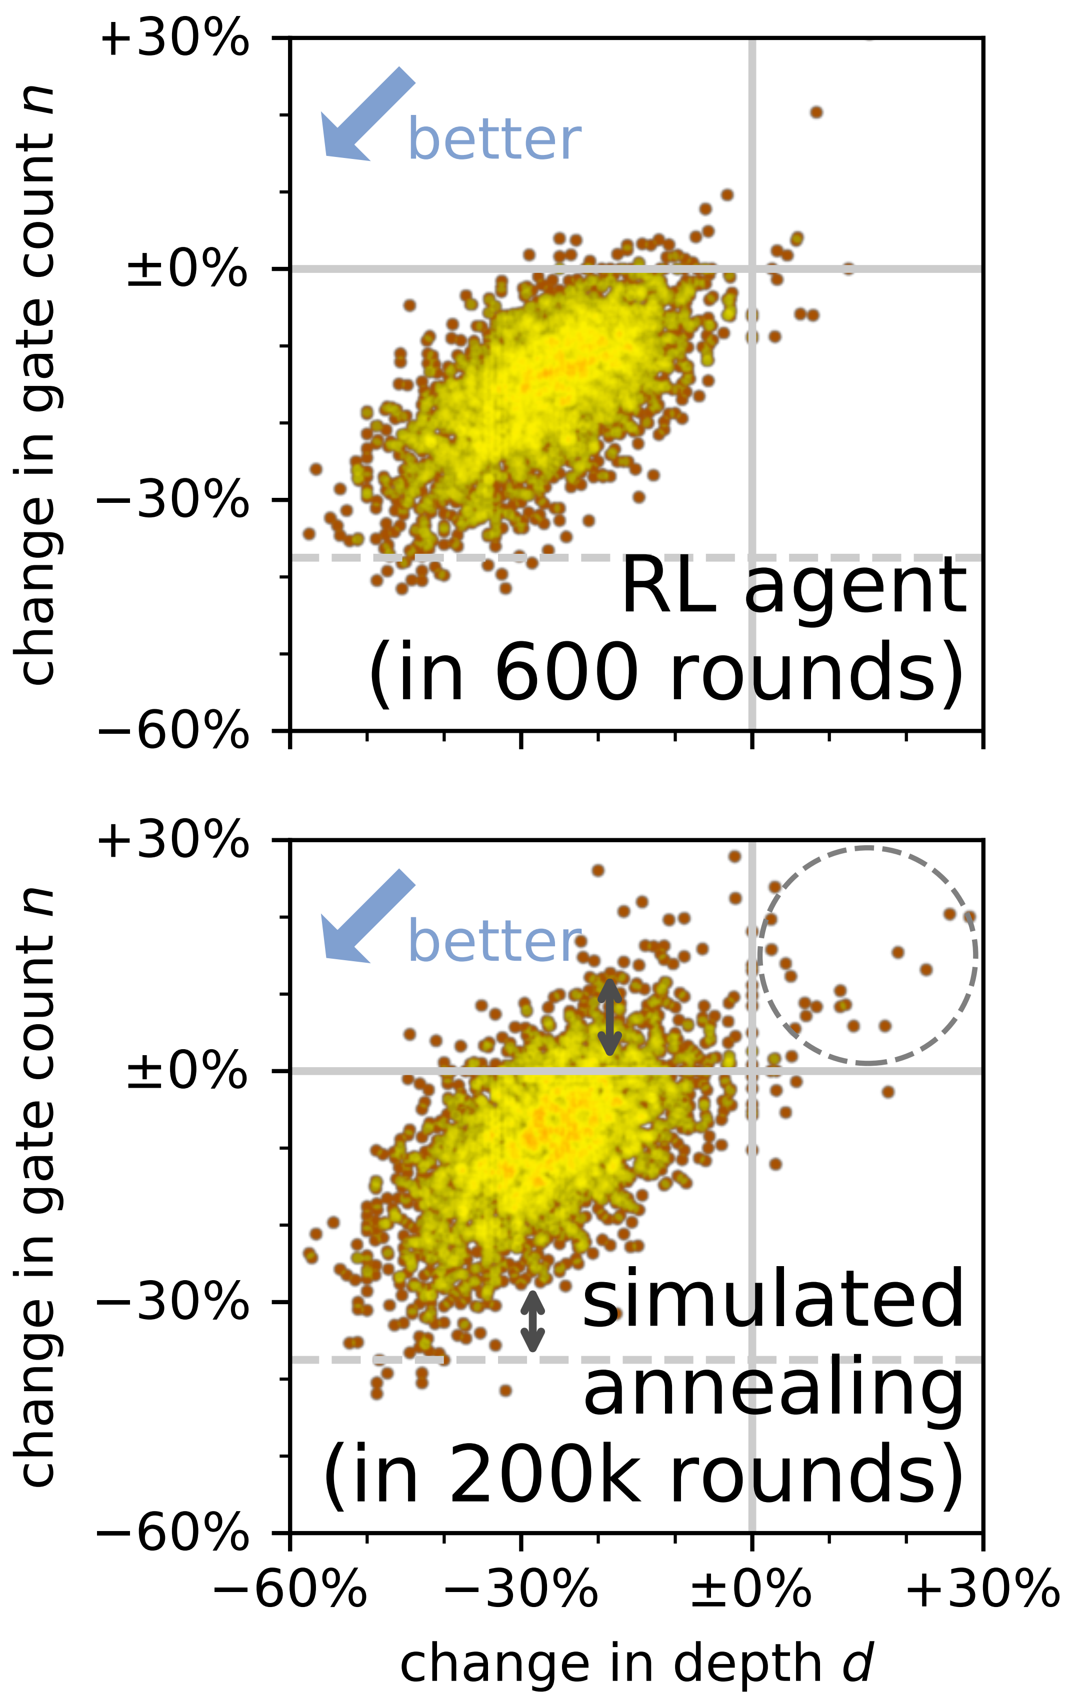

RL (after 1-week training):
2 min per circuit

works for arbitrary circuits!

Simulated annealing:
1-3 d per circuit

Large-scale Random Circuits
(same agent, now applied to larger circuit)

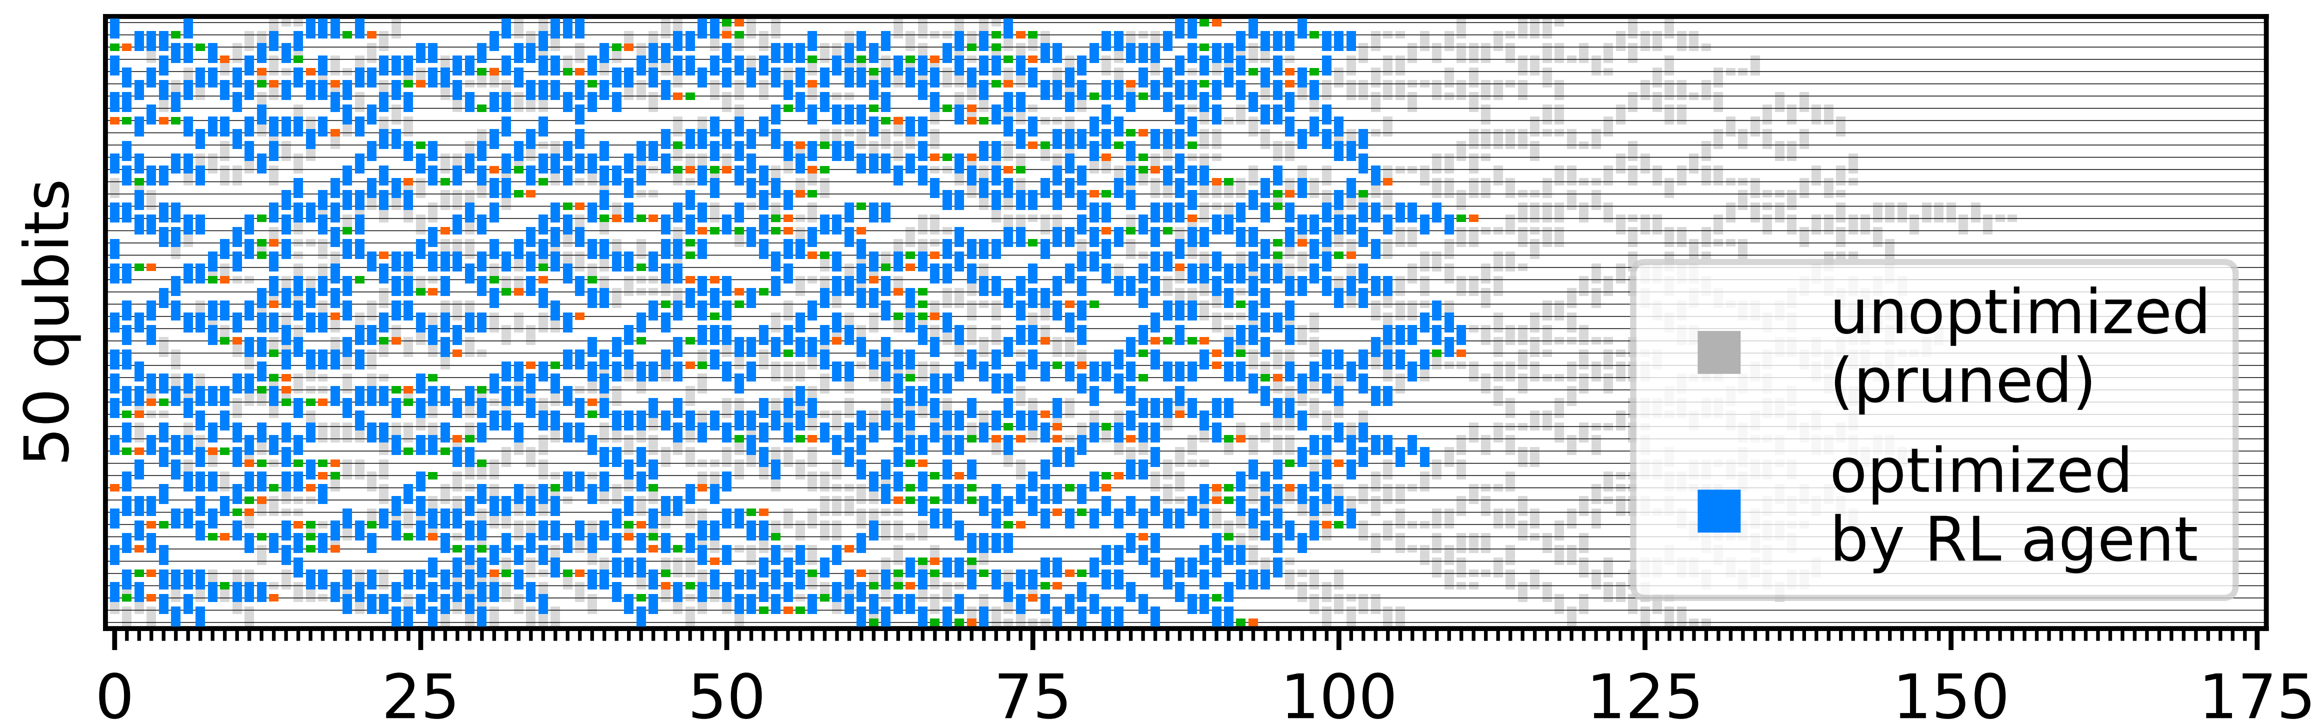

Convolutional network permits successful transfer of learned behaviour to much larger circuits: local environment of gates is relevant!

simulated annealing: ~ 1 week, comparable to full training time for general RL agent (that runs in 3-5 h)

Application to a real algorithm

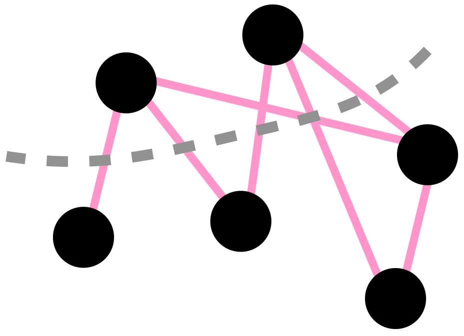

Example:
Quantum Approximate Optimization Algorithm (QAOA)
– specifically, for the MaxCut problem

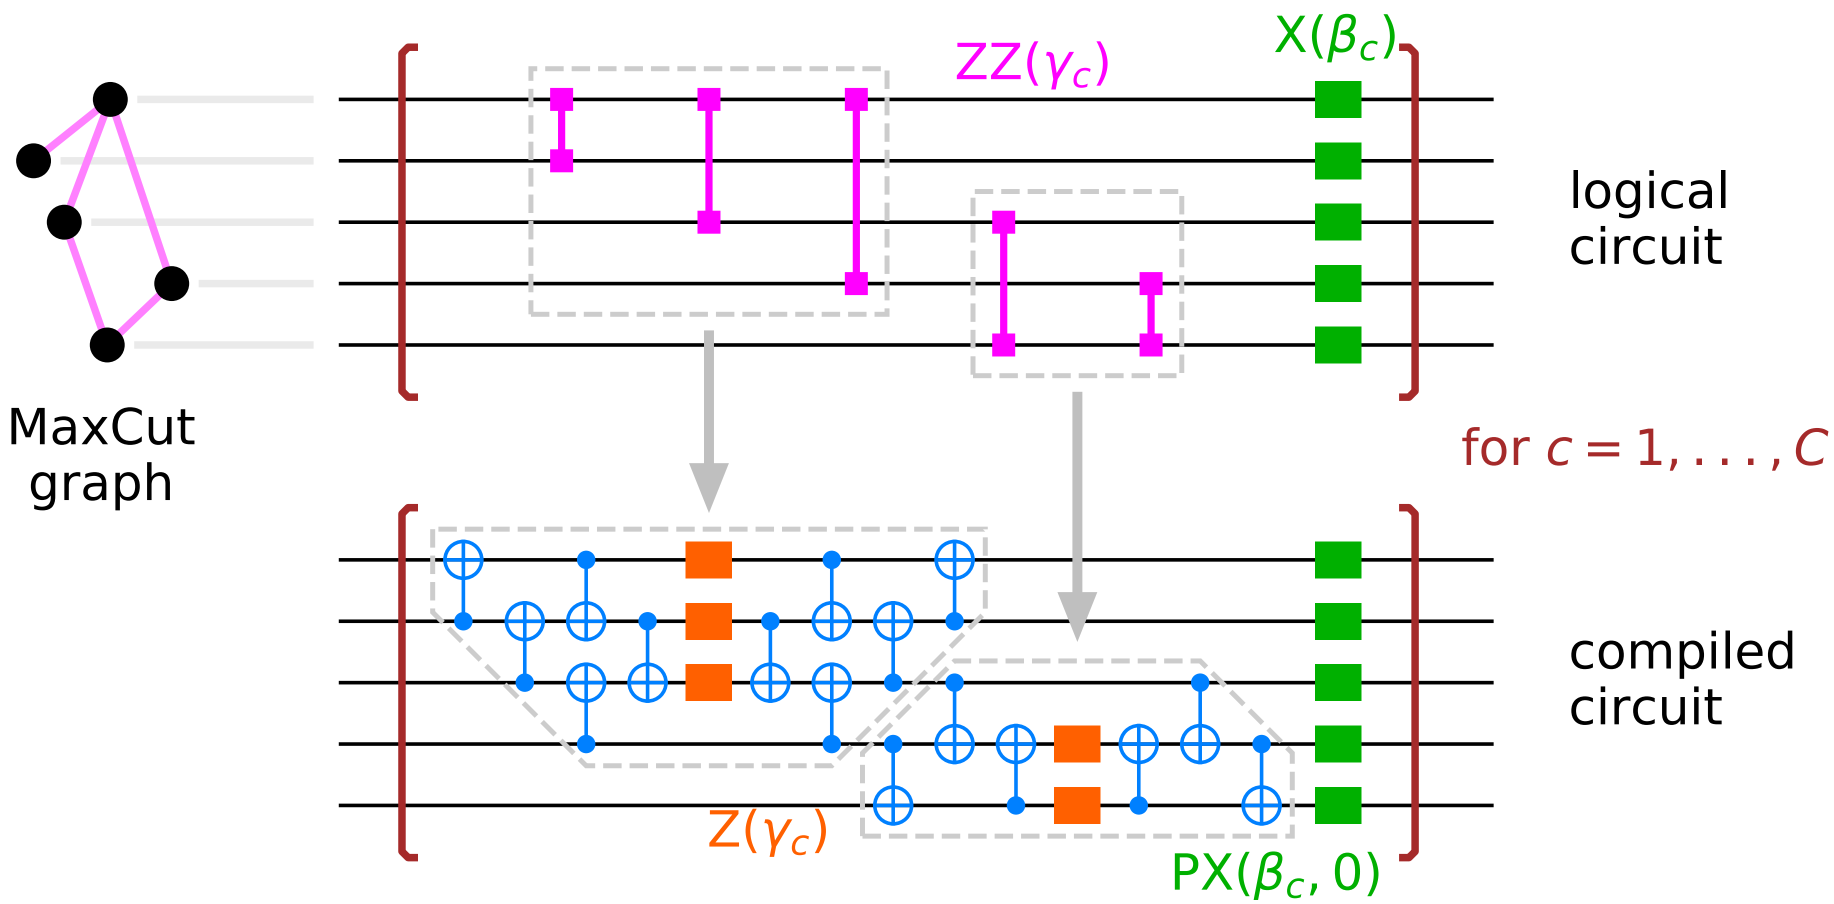

QAOA: Farhi et al, 2014
Experimental MaxCut-QAOA (Google): Harrigan et al, 2021

MaxCut Circuit Optimization

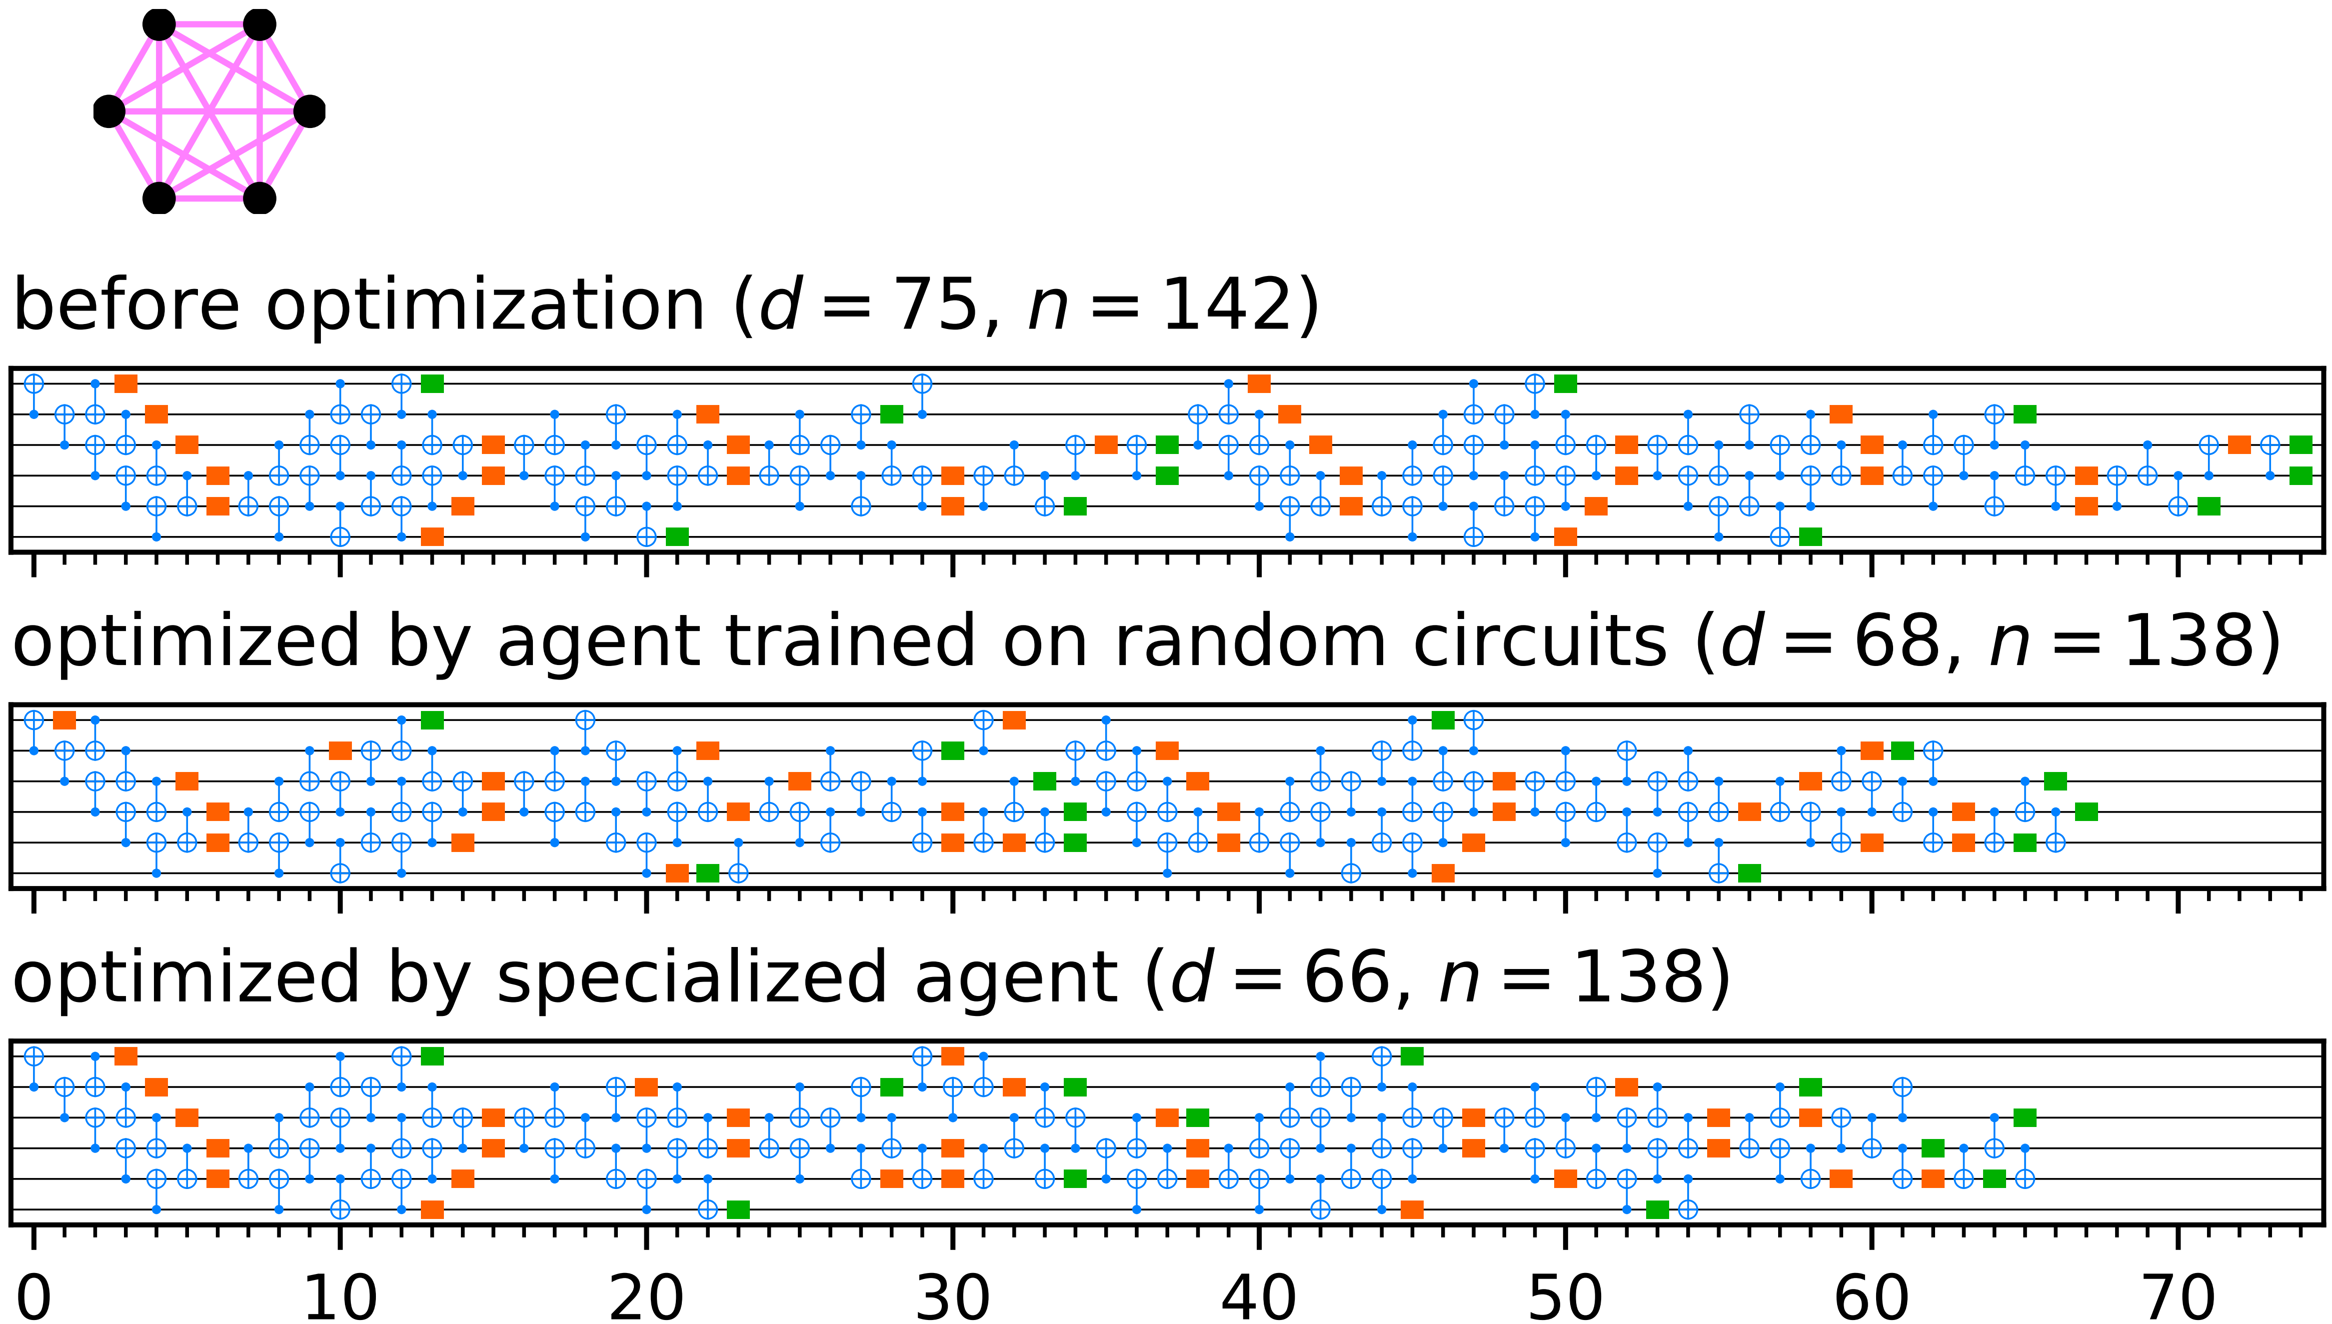

Future: Quantum Circuit Dataset

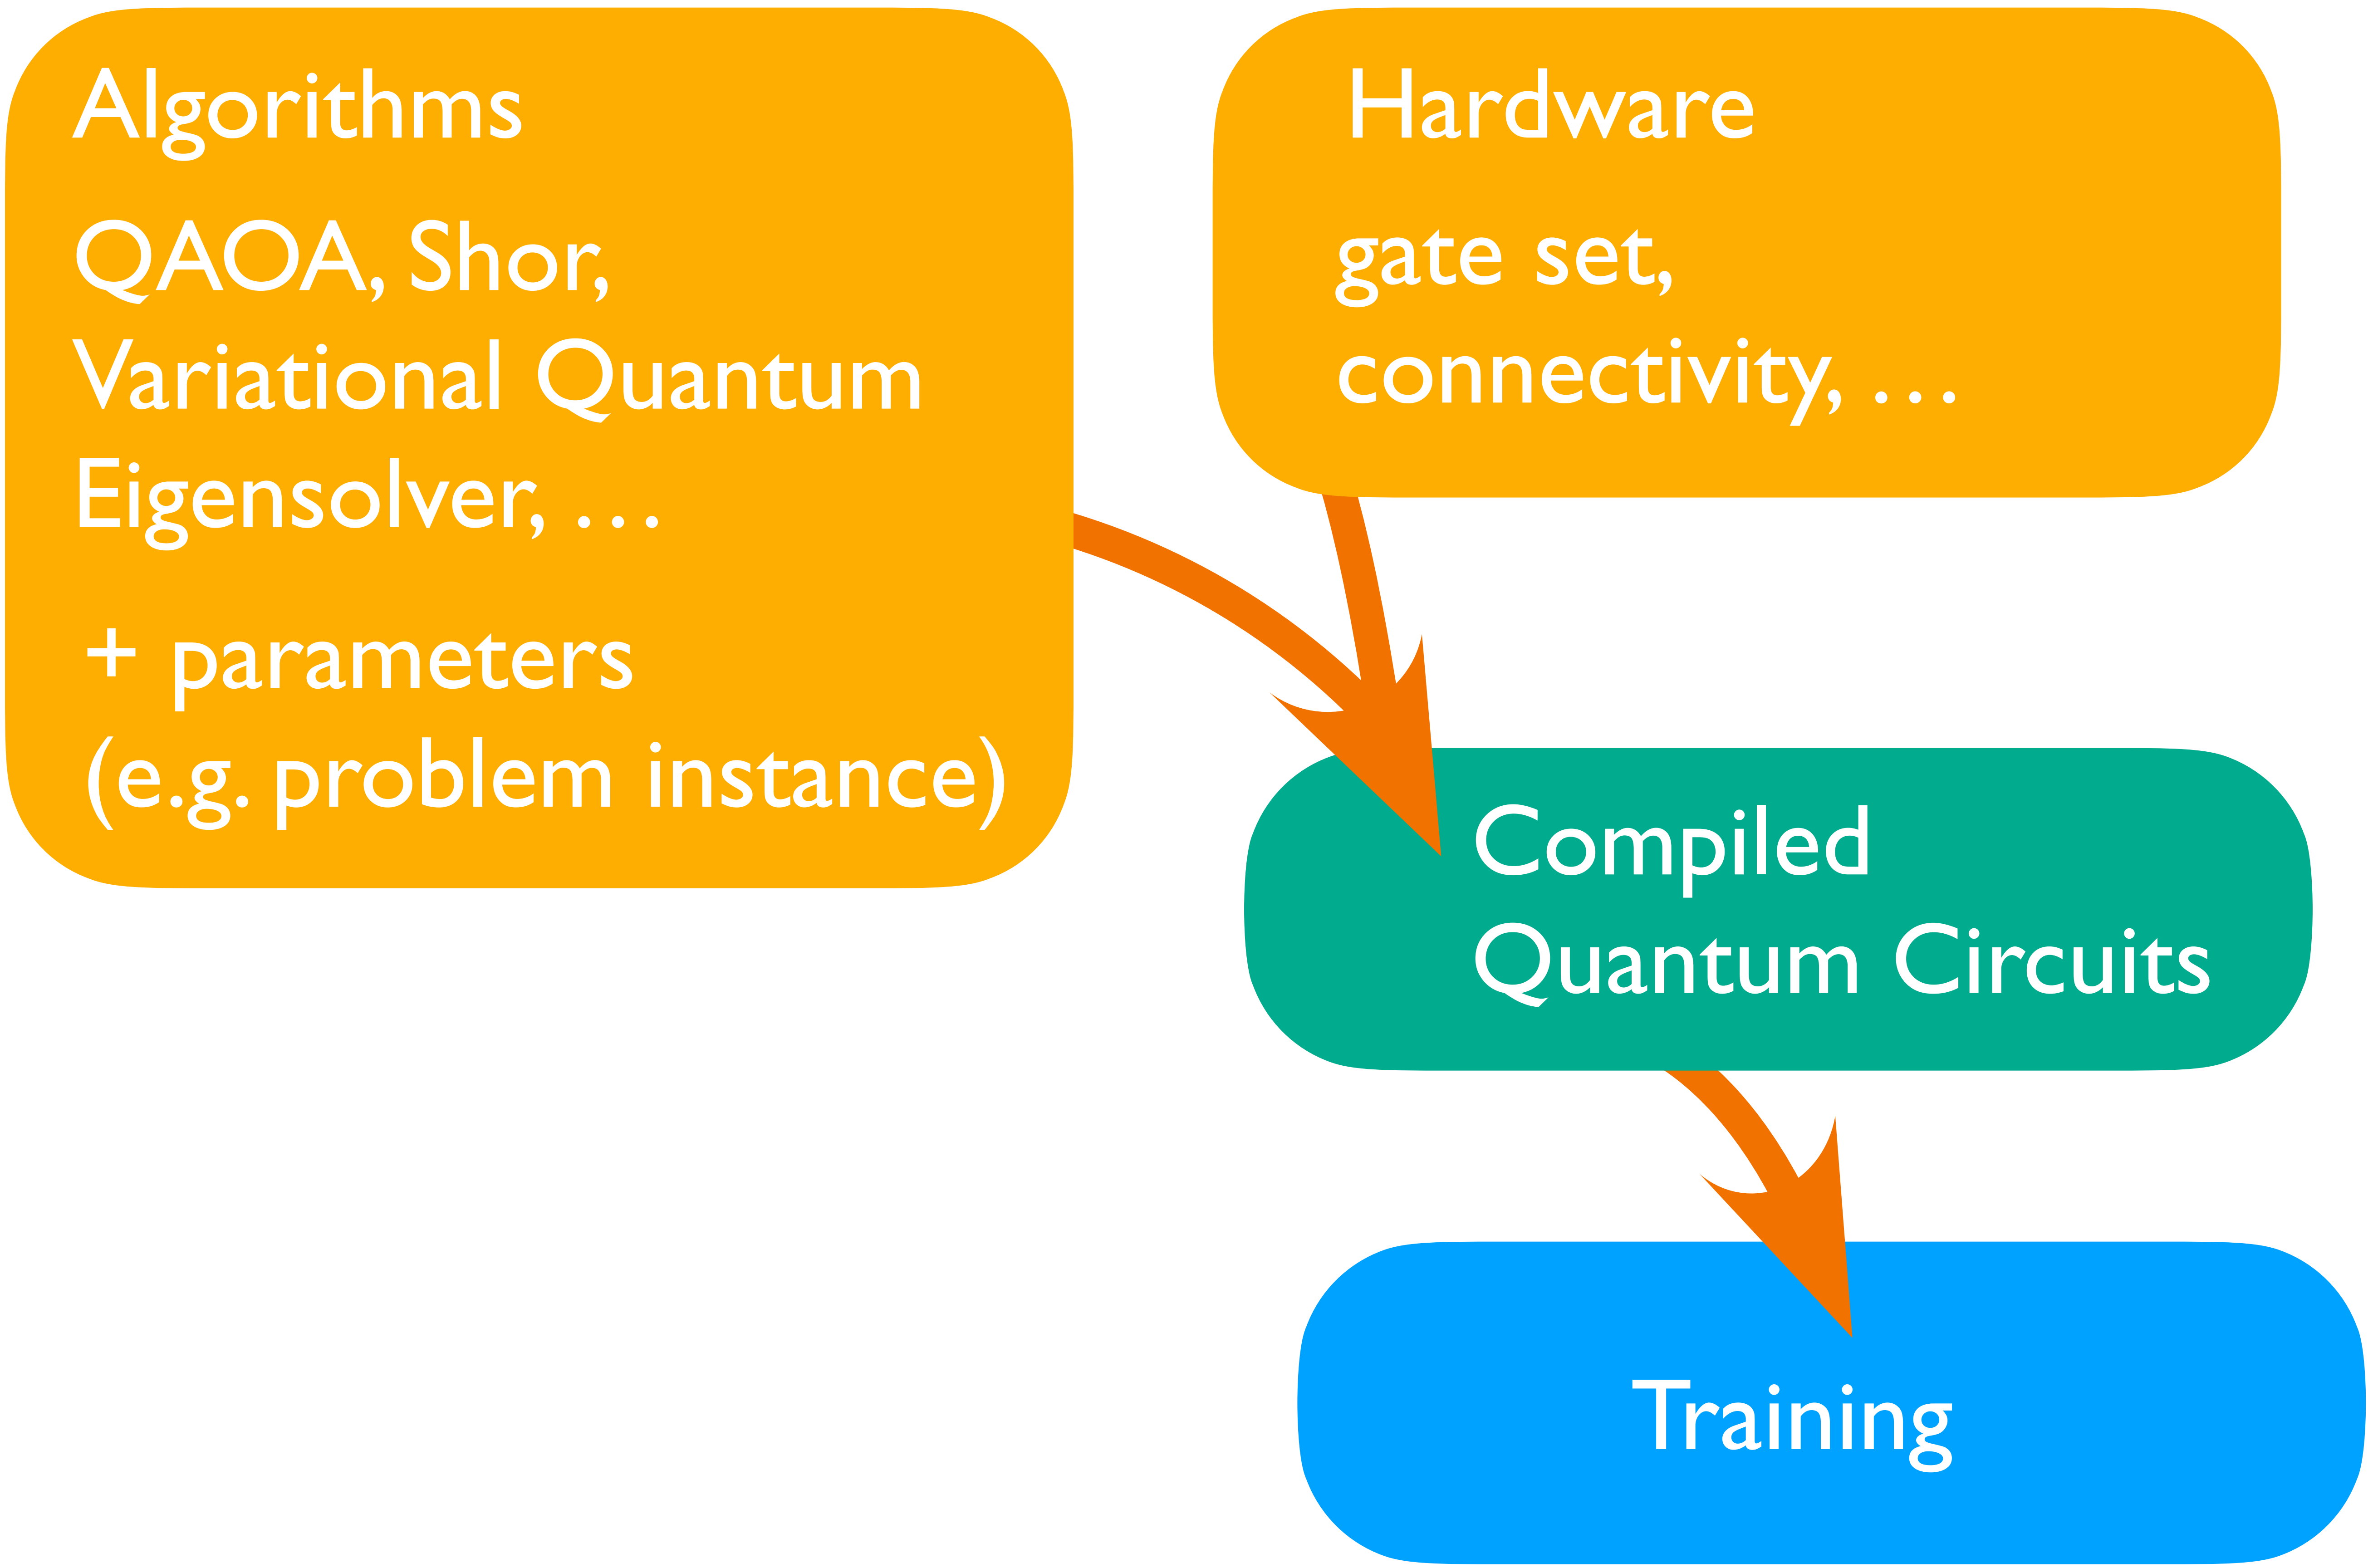

Lecture Notes "Machine Learning for Quantum Devices"

arXiv 2101.01759 - SciPost Lecture Notes 2021

Current online lecture series "Advanced Machine Learning for Physics, Science, and Artificial Scientific Discovery"

Discovering Quantum Error Correction Strategies

work with **Thomas Fösel, Petru Tighineanu, and Talitha Weiss**

Physical Review X 031084 (2018)

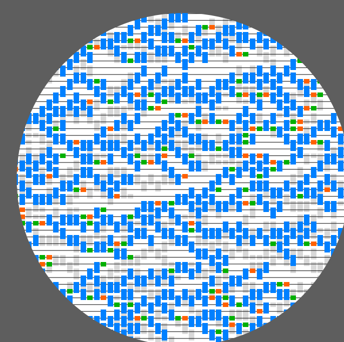

Quantum Circuit Optimization

work with **Thomas Fösel, Murphy Yuezhen Niu, and Li Li**

arXiv 2103.07585

New: Several Postdoc positions available in Reinforcement Learning applied to Quantum Technologies

- "Munich Quantum Valley" Initiative

(ask me, Florian.Marquardt@mpl.mpg.de; deadline Nov 15)