Title: The Quantum Cartpole

Speakers: Evert van Nieuwenburg

Collection: Machine Learning for Quantum Many-Body Systems

Date: June 16, 2023 - 10:00 AM

URL: https://pirsa.org/23060048

Abstract: How do you control something you can not look at? For controlling quantum systems, information on the system's state could come through weak measurements. Such measurements provide some information, but will inevitably also perturb the system, meaning there is noise both in the state estimation as well as in the measurement. We study a simple single particle quantum setup (the quantum equivalent of the instability problem known as the cartpole problem) and investigate several control methods including reinforcement learning, and compare their performance.

The Quantum Cartpole

Evert van Nieuwenburg
Applied Quantum Algorithms @ Leiden University

Leiden Institute of Physics

liacs

Not So (?)

The Quantum Cartpole

Evert van Nieuwenburg

Applied Quantum Algorithms @ Leiden University

Leiden Institute of Physics

liacs

The aQa group covers the full quantum pipeline

Vedran Dunjko

Jordi Tura

Hao Wang

Alfons Laarman

Me 😅

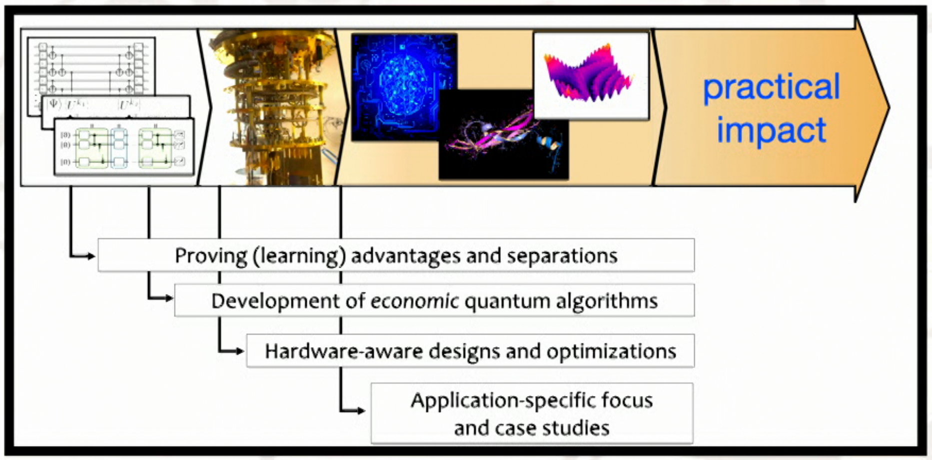

- **Circuit cutting bounds**
- **Hybrid classical-quantum for tree-search**
- **Learning separation classical and QML**
- **Quantum adv. beyond kernel methods**
- **...**

Pioneered at aQa

Please reach out for opportunities

Our focus is rather more ‘aaq' than ‘aqa’

Algorithms applied to Quantum **Applied Quantum Algorithms**

De-noising of STM data to identify gaps

Quantum error decoding using graph neural networks

Genetic algorithms for optimising average gradients in quantum circuits

Felix Frohnert

Representation Learning of Quantum Systems

Representation Learning of small quantum circuits

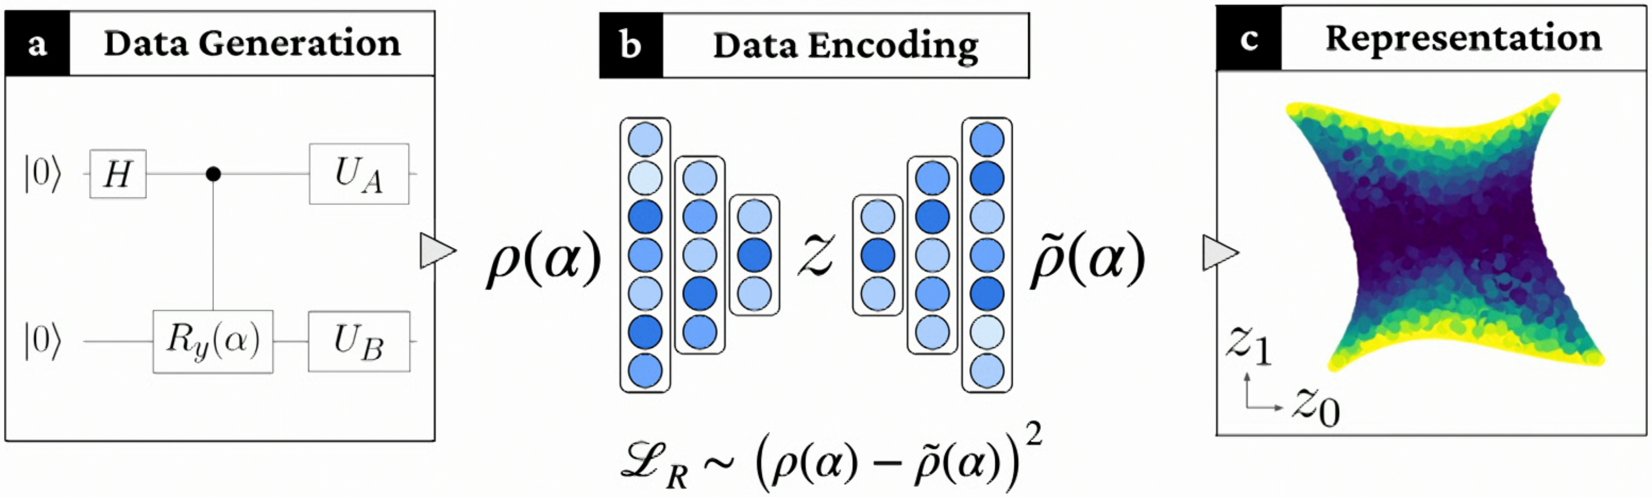

https://arxiv.org/abs/2306.05694

1/3

Representation Learning of small quantum circuits

$$\mathcal{L}(\boldsymbol{x};\boldsymbol{\phi},\boldsymbol{\theta}) = \mathcal{L}_R(\boldsymbol{x};\boldsymbol{\phi},\boldsymbol{\theta}) + \beta \cdot \mathcal{L}_{KL}(\boldsymbol{x};\boldsymbol{\phi},\boldsymbol{\theta})$$

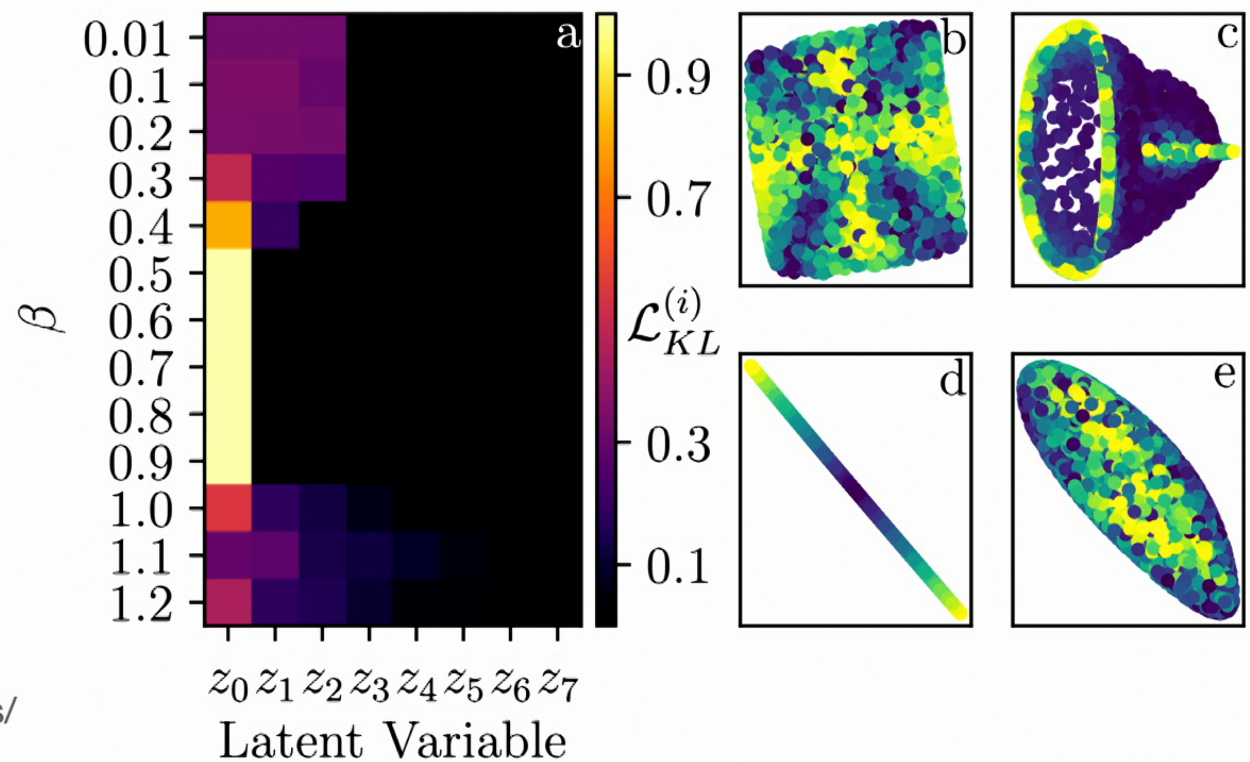

https://arxiv.org/abs/2306.05694

2/3

Representation Learning of small quantum circuits

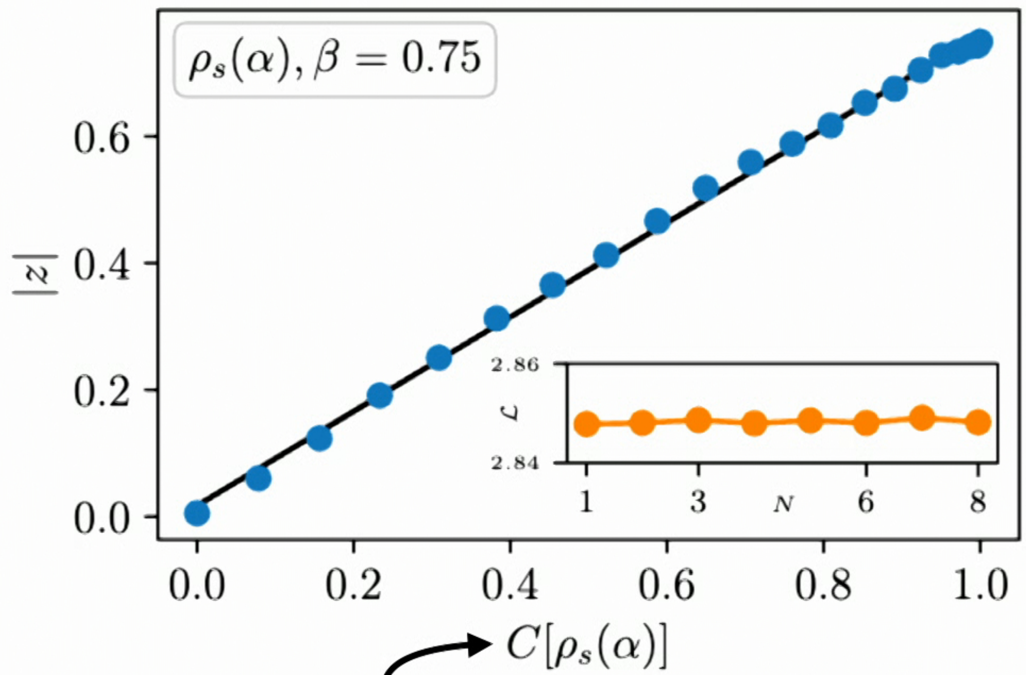

Concurrence (mixed state entanglement measure)

https://arxiv.org/abs/2306.05694

3/3

Not So (?)

The Quantum Cartpole

Evert van Nieuwenburg

Applied Quantum Algorithms @ Leiden University

Leiden Institute of Physics

liacs

(I think) you should care about quantum games

1) Develop new talent, raise awareness

~40% of global population plays video games!*

2) Drive hardware

Why do you buy a new phone?

3) Fruitful research framework

Game circuits as benchmark? AI for quantum game?

*https://explodingtopics.com/blog/number-of-gamers

beta.tiqtaqtoe.com

More info?

evert@tiqtaqtoe.com

Powered by
Quantum
Delta NL

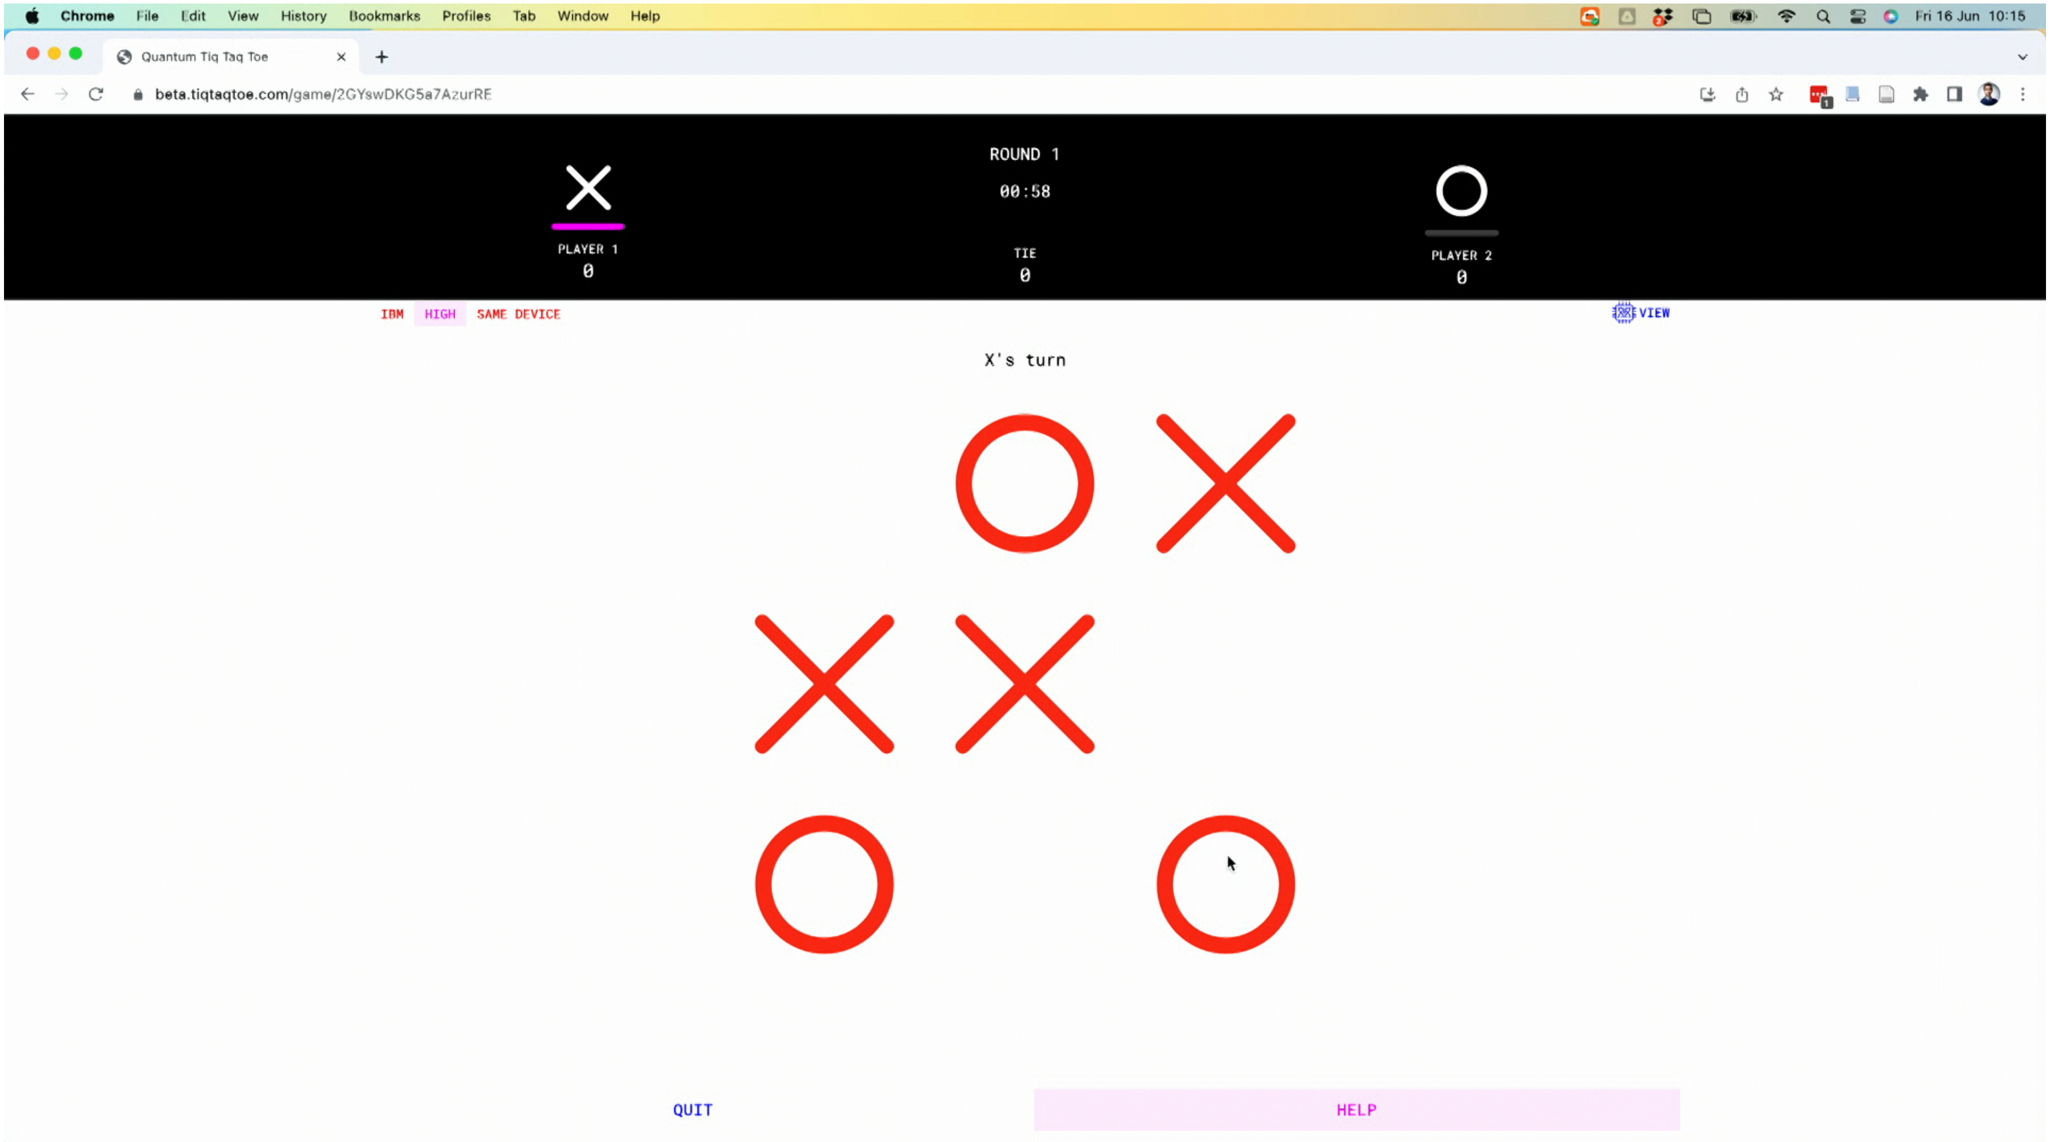
Chrome File Edit View History Bookmarks Profiles Tab Window Help
Fri 16 Jun 10:15
Quantum Tiq Taq Toe
beta.tiqtaqtoe.com/game/2GYswDKG5a7AzurRE
ROUND 1
00:58
PLAYER 1
0
TIE
0
PLAYER 2
0
IBM
HIGH
SAME DEVICE
VIEW
X's turn
QUIT
HELP

The Quantum Cartpole

Or: how do you control something you cannot look at?

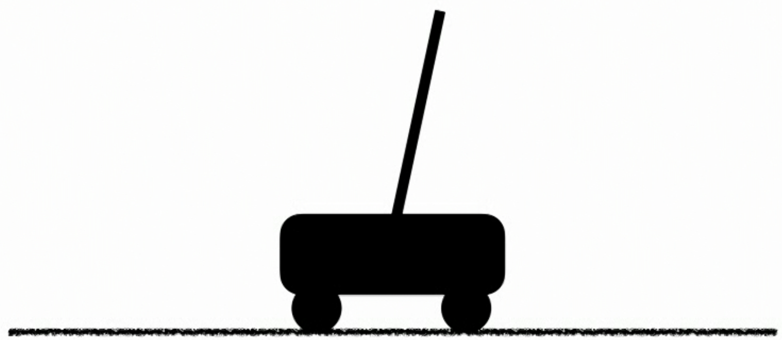

Kai Meinerz

Simon Trebst

Mark Rudner

The classical cartpole is a standard benchmark

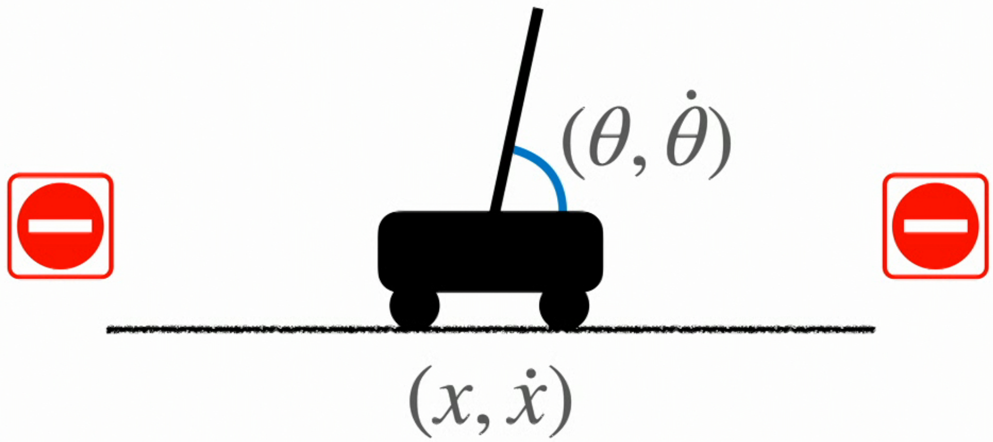

1. At every time step, you get $s = (x, \dot{x}, \theta, \dot{\theta})$

2. You apply a force to the cart

3. You continue for as long as
1) the pole doesn't fall over
2) the cart is within bounds

The fundamental problem is the same as this instability

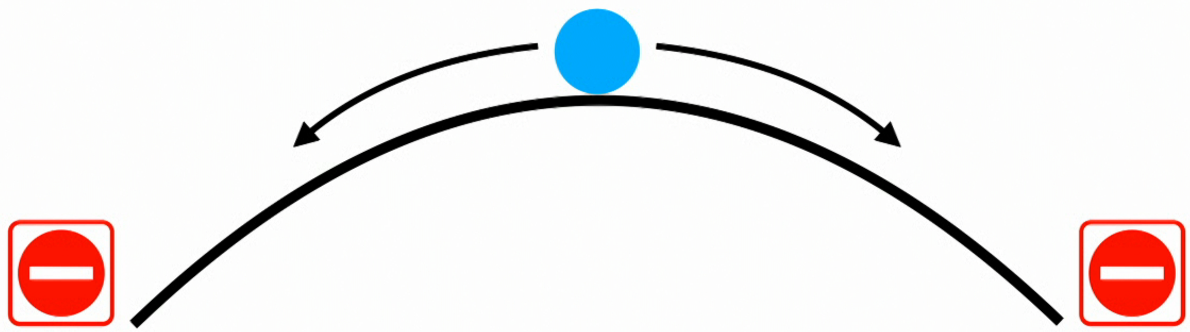

1. At every time step, you get $s = (x, \dot{x})$

2. You apply a force to the ball

3. You continue for as long as
1) the ball doesn't go out of bounds

The classical version can be optimally controlled

…under some assumptions…

Assume small deviations -> linearize

$$s_{t+1} = As_t$$

Enter LQR (Linear Quadratic Regulator)

$$s_{t+1} = As_t + Bu_t$$

the input/control

$$u_t = -Ks_t$$

Determine input s.t. $x = 0$, $\dot{x} = 0$

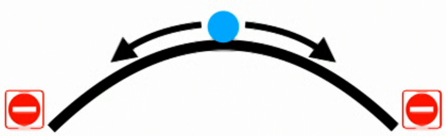

$x = (x, \dot{x})$

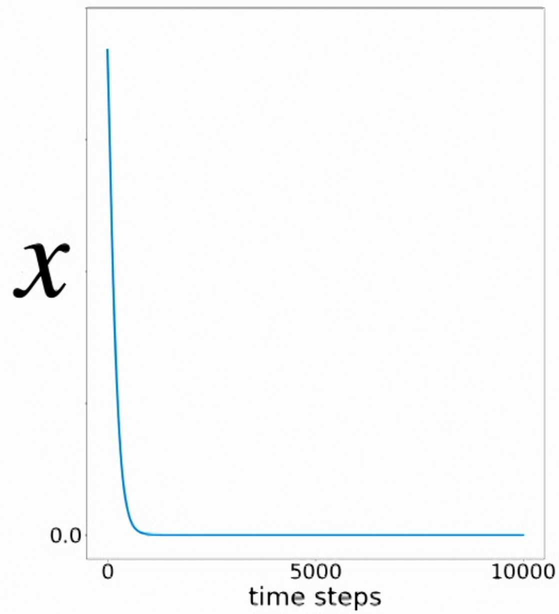

It can still be optimally controlled if we make it a little harder

At every time step, you get $s = (x)$ (i.e. stationary picture)

~~$s_{t+1} = As_t + Bu_t$ with $u_t = -Ks_t$~~

Kalman Filter aka **L**inear **Q**uadratic **E**stimator

$$s_t = Fs_{t-1} + Bu_t$$

LQE looks back in time and estimates s_t Is used by LQR to control the system

Even adding noise is not a (big) issue

…under some assumptions…

$$s_{t+1} = As_t + Bu_t + w_t$$
$$y_t = Cs_t + v_t$$

But requires knowing (w_t, v_t)

$$w_t = \mathcal{N}(0, W_t)$$

Gaussian white noise process

Linear Quadratic Gaussian Controller

LQE looks back in time and estimates s_t Is used by LQR to control the system

Even adding noise is not a (big) issue

…under some assumptions…

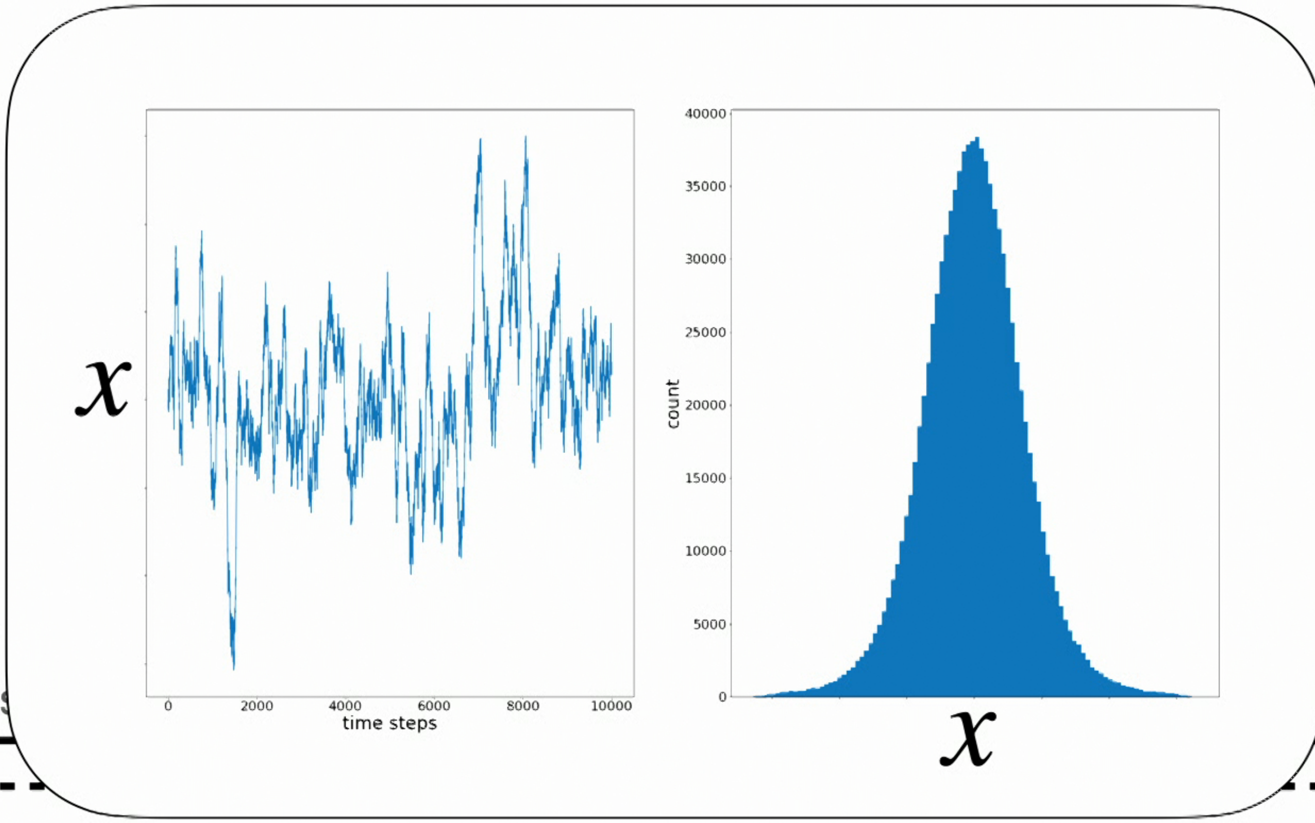

LQE looks

e system

Perhaps more measurements = better state estimate?

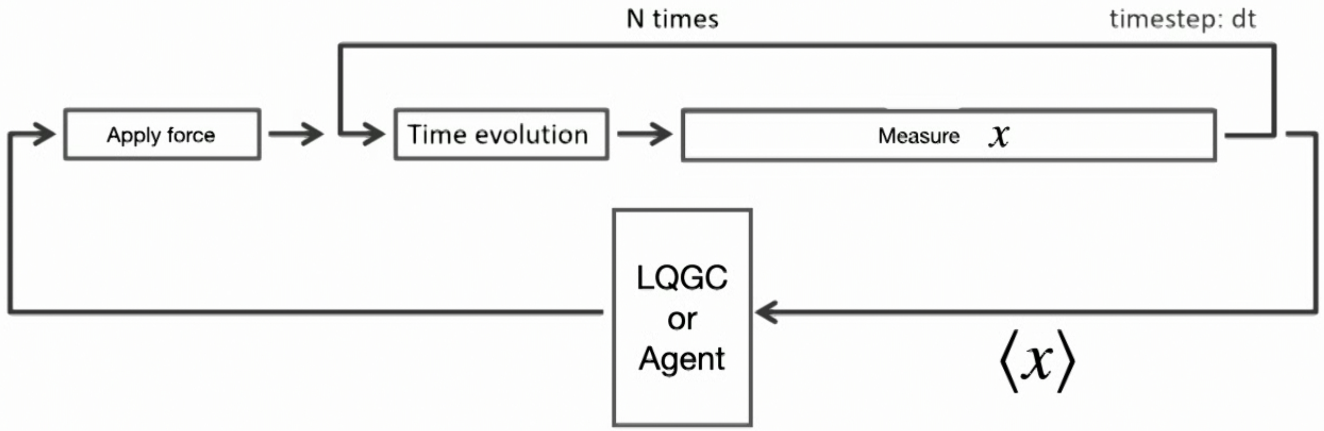

1. At every step, you get N measurements

2. You apply a force

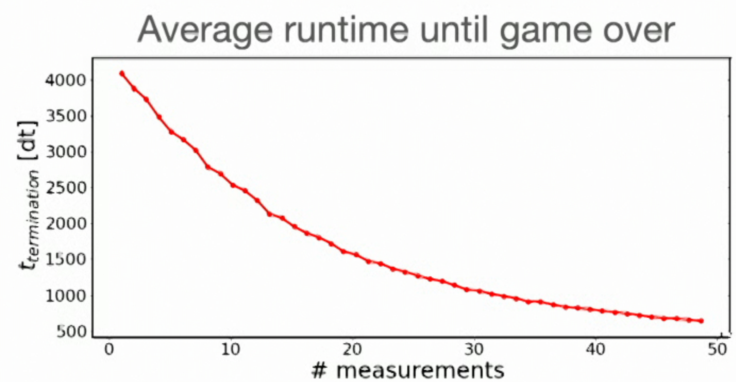

Non-linearity is not a (big) problem either

Extended Kalman Filter

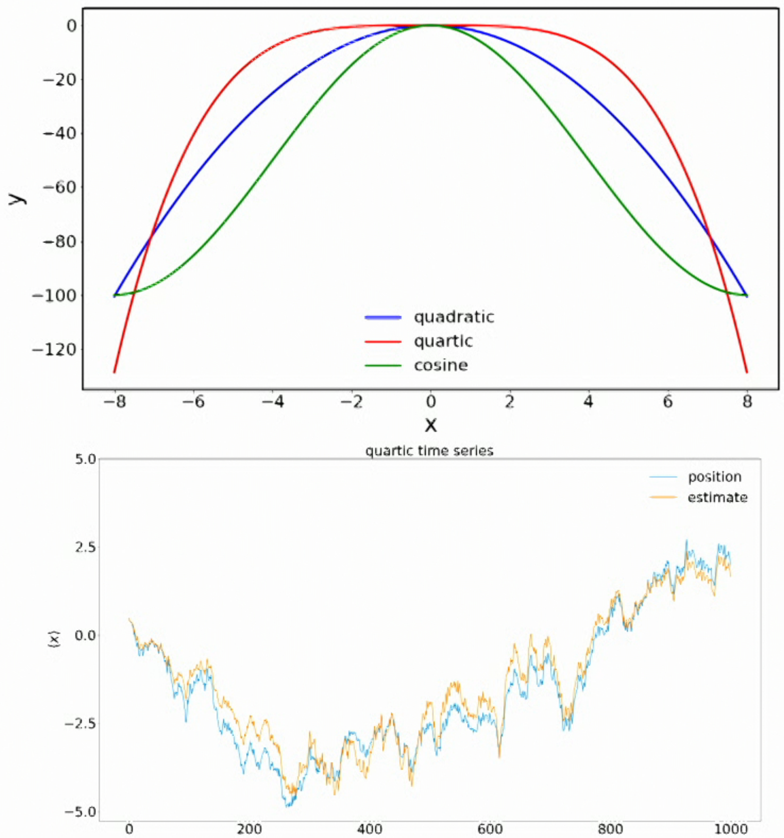

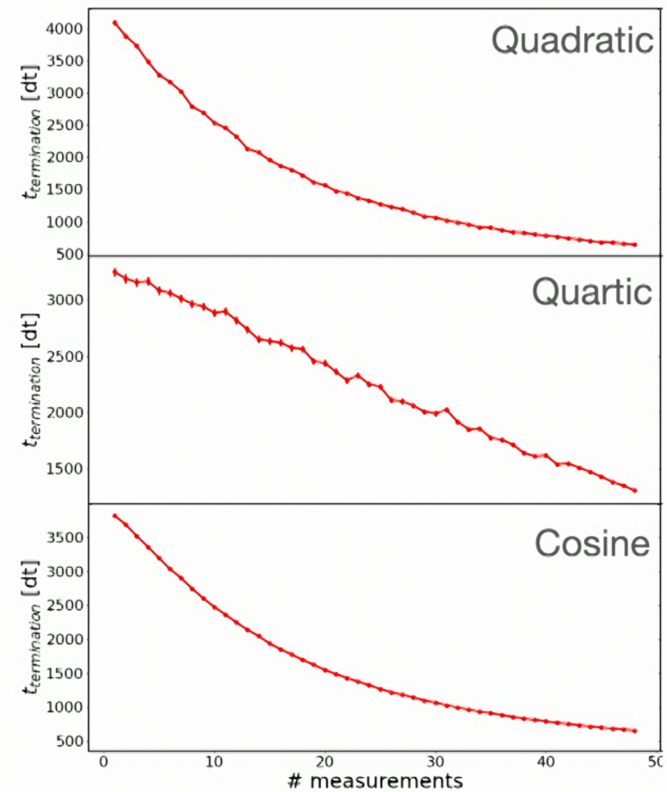

Final resort: let's make it quantum

Deep Reinforcement Learning Control of Quantum Cartpoles

Zhikang T. Wang,[1,*] Yuto Ashida,[2] and Masahito Ueda[1,3]

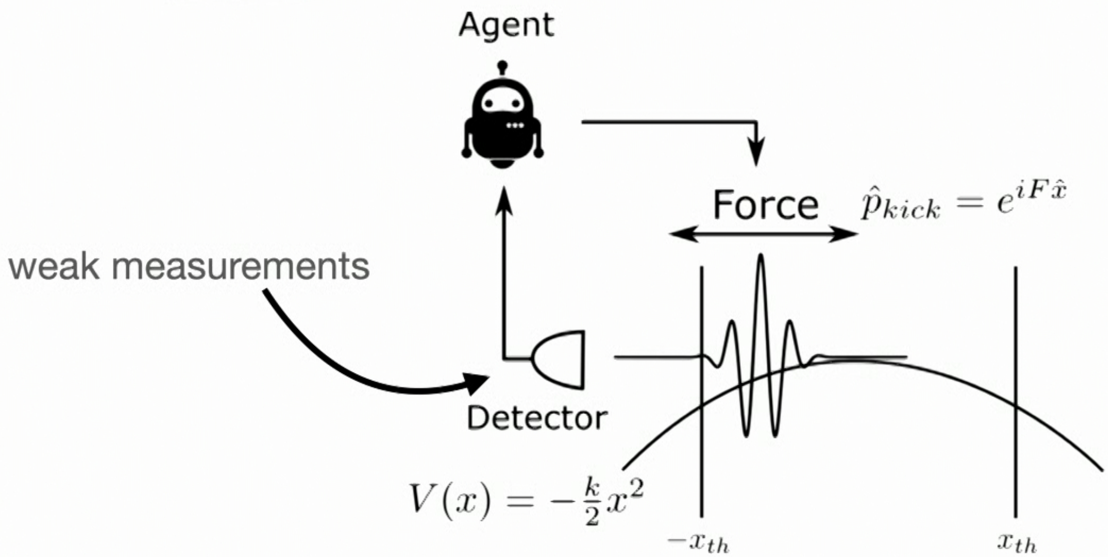

* Algorithms such as GRAPE etc are gradient-based, and work for isolated non-stochastic systems.

Now using weak measurements to get estimates

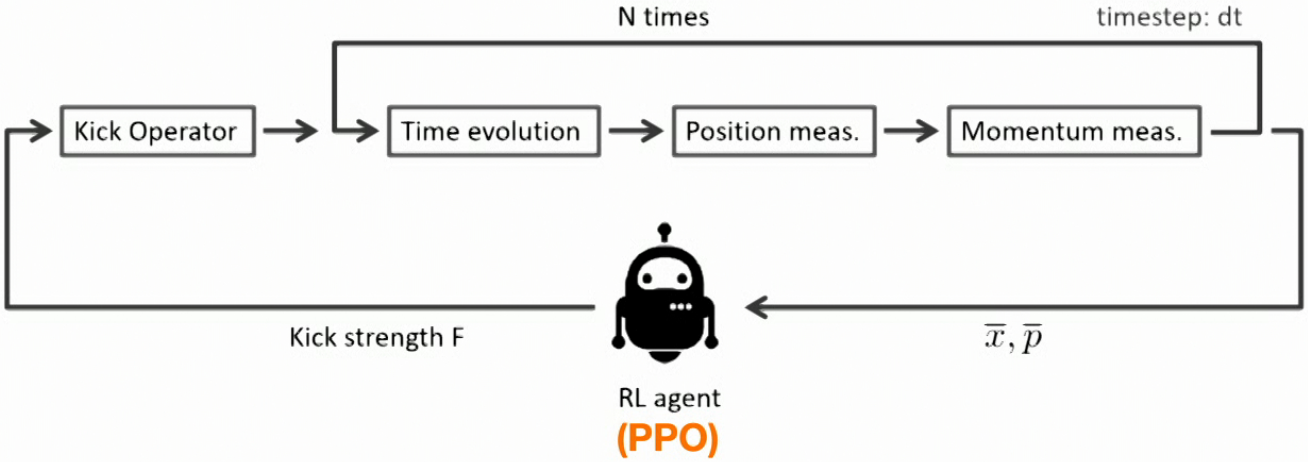

1. At every step, you get 2N weak measurements

2. You apply a force to the particle

The RL agent learns to control the quantum cartpole

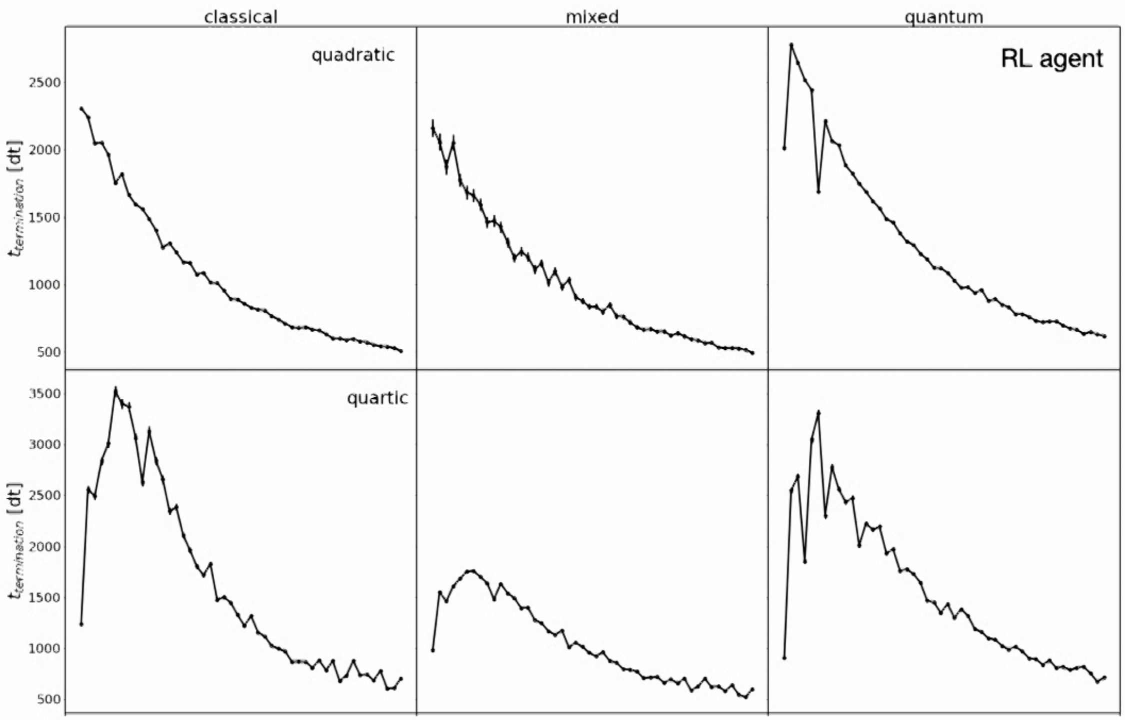

Mixed
=
RL trained on classical, evaluated on quantum

The RL agent learns to control the quantum cartpole

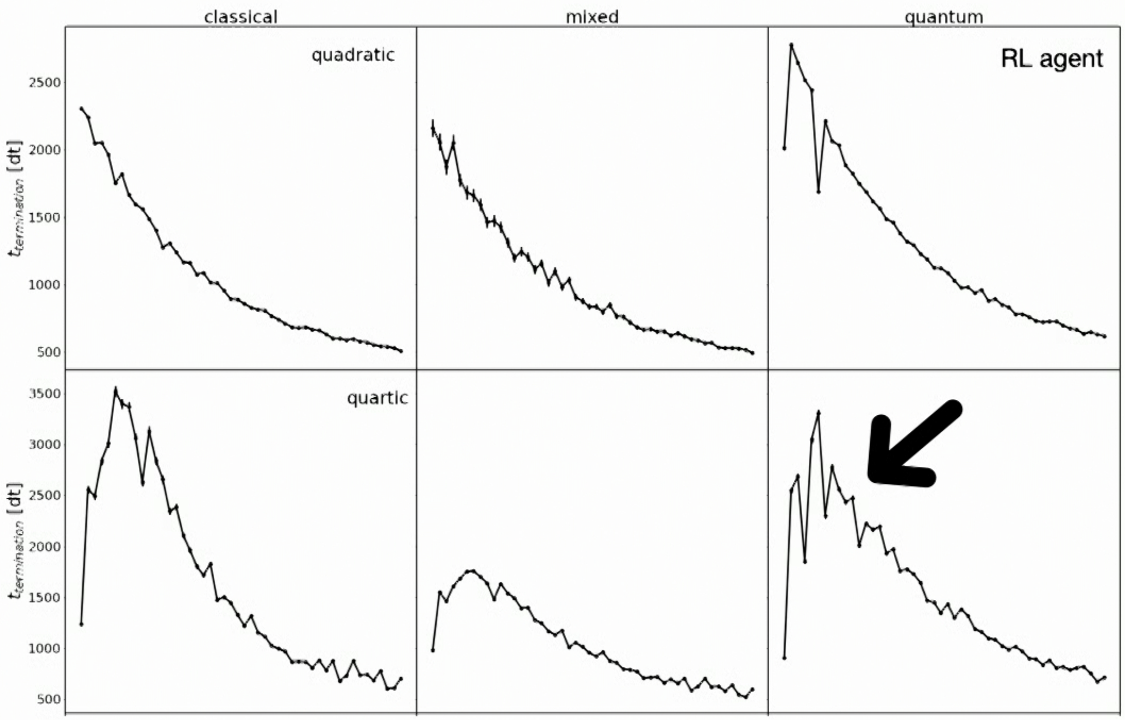

Mixed
=
RL trained on classical, evaluated on quantum

RL can outperform the classical standard

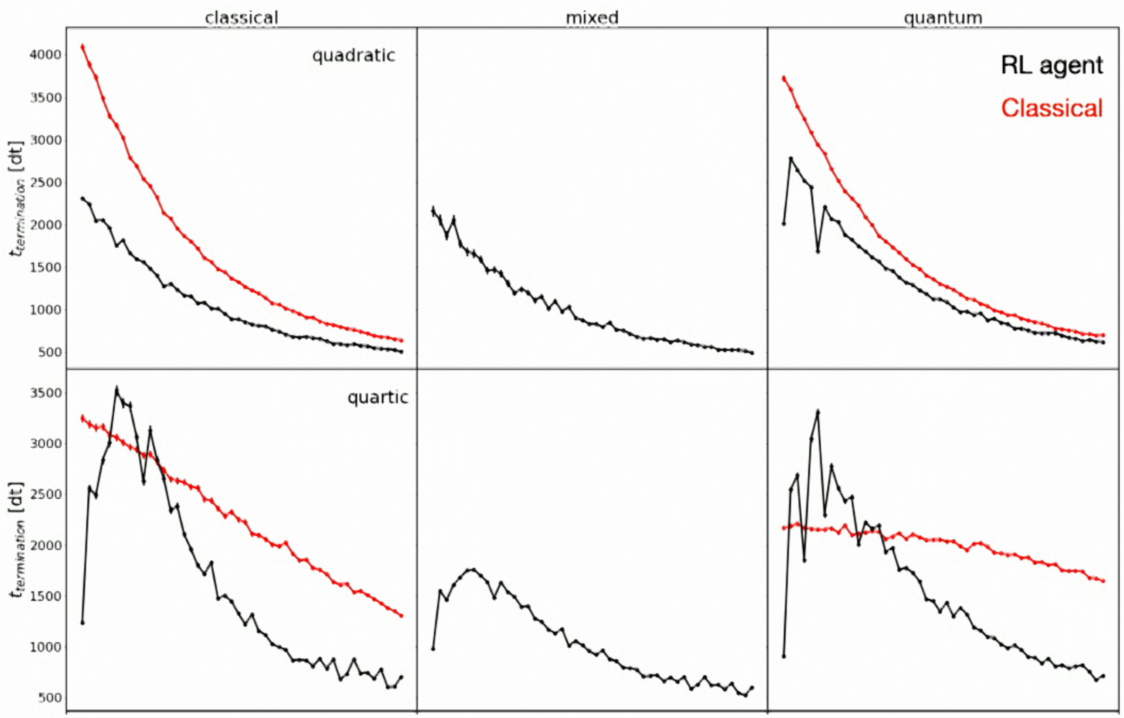

Mixed
=
RL trained on classical, evaluated on quantum

The Quantum Cartpole

Concluding

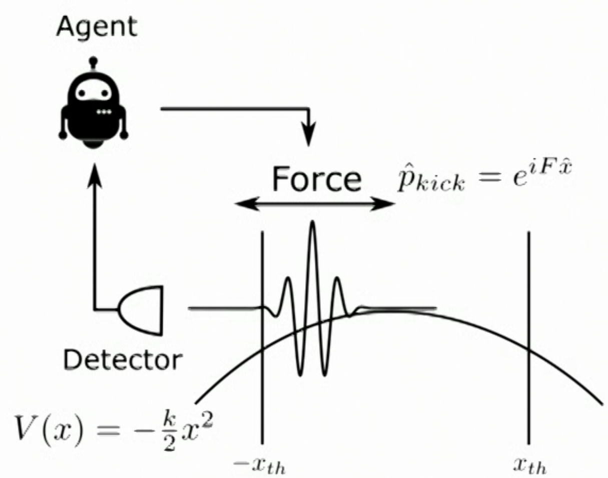

1. A classical stochastic controller can control quantum systems using weak measurements
2. For non-linear (and noisy) cases, RL controller is able to do better

Thank you!