Machine learning methods for Robust Quantum Optimal Control

Irtaza Khalid

Road Map

1. Backdrop
2. Robustness certification
    a. Optimality is not enough
    b. Connection with classical robust control
3. Optimisation of quantum dynamics
    a. Reinforcement learning without a model
    b. Model learning for control

1. Backdrop

Quantum tech offers certain fundamental advantages

Communications (Bennett and Brassard, BB84 protocol, 1984)

Sensing (spin squeezing Wineland et. al. 1992)

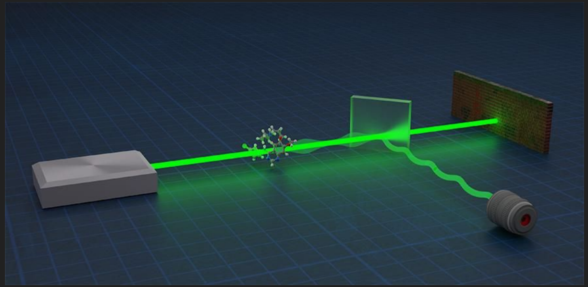

Boson sampling (Arute et. al. 2019)

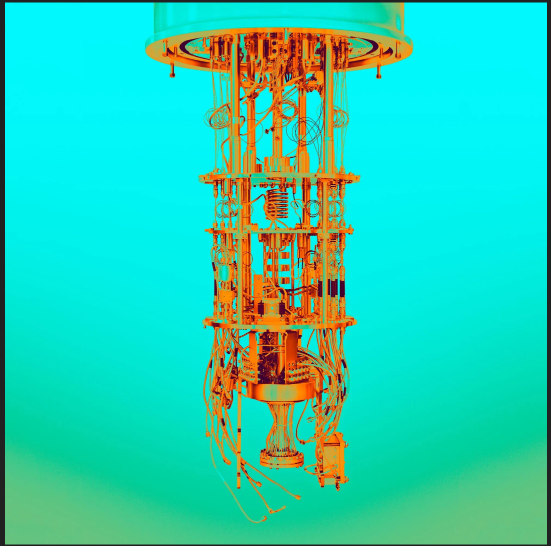

Simulation (cold atoms in optical lattice Jaksch et. al. 1998)

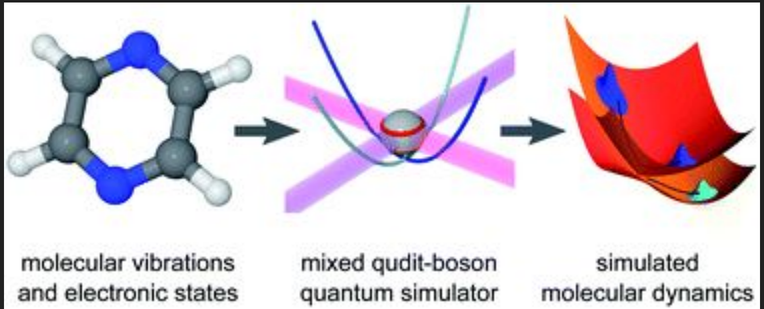

Cryptography (Shor et. al.1994)

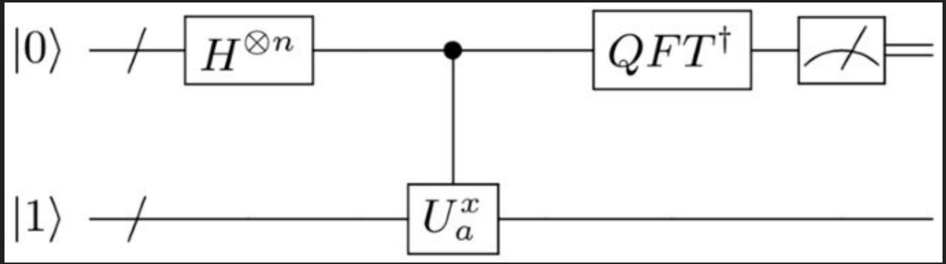

Noise is the biggest bottleneck between idea and reality

1. Decoherence / loss of “quantumness” due to environmental interactions
2. Unitary noise, e.g. model error
3. Measurement / estimation noise

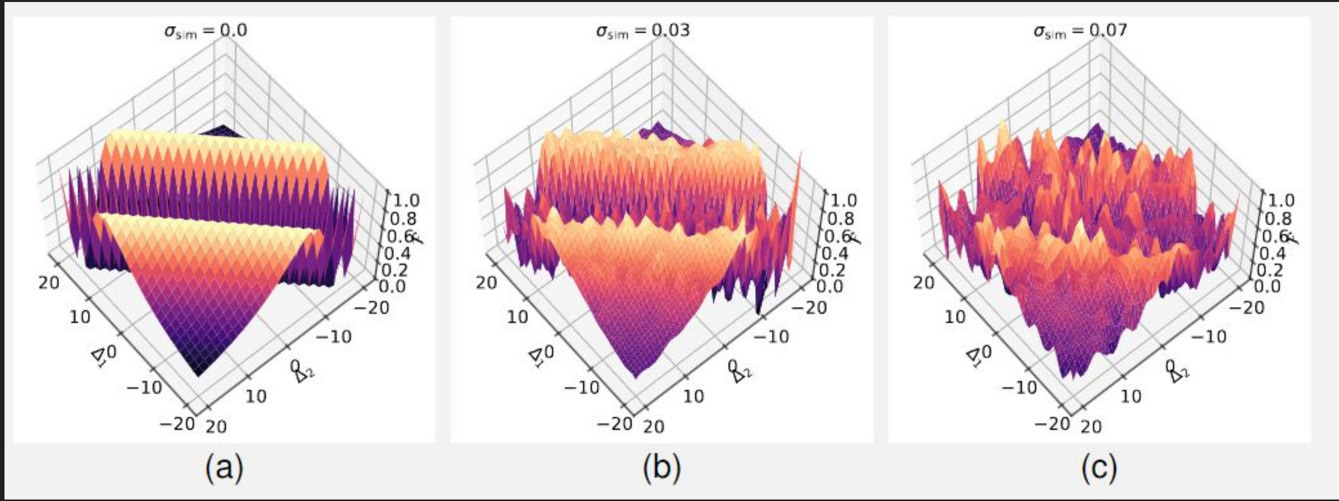

Towards robust optimal data-driven control

1. Optimal control consists of a system and a controller (also part of the system that is variable) $$H(\mathbf{u}(t), t) = H_0 + H_c(\mathbf{u}(t), t)$$

2. The evolution is under some version of the Schrödinger DE
$$\frac{d}{dt}U(\mathbf{u}(t), t) = -\frac{i}{\hbar}H(\mathbf{u}(t), t)U(\mathbf{u}(t), t), \quad U(t = t_0) = \mathrm{I}$$

3. The figure-of-merit or target that is optimised is:
$$\mathcal{F}(U_{\text{target}}, U(t_0, t_1, \mathbf{u}(t))) = \frac{1}{2^{2n}} \left| \mathrm{Tr}\left[U^\dagger_{\text{target}} U(t_0, t_1, \mathbf{u}(t))\right] \right|^2$$

Towards robust optimal data-driven control

$$H(\mathbf{u}(t), t) = H_0 + H_c(\mathbf{u}(t), t)$$ + noise

$$\frac{d}{dt}U(\mathbf{u}(t), t) = -\frac{i}{\hbar}H(\mathbf{u}(t), t)U(\mathbf{u}(t), t), \quad U(t = t_0) = \mathrm{I}$$

+ environmental decay

$$\mathcal{F}(U_{\text{target}}, U(t_0, t_1, \mathbf{u}(t))) = \frac{1}{2^{2n}} \left| \mathrm{Tr}\, [U^{\dagger}_{\text{target}} U(t_0, t_1, \mathbf{u}(t))] \right|^2$$

Is its optimal value stable under perturbations?

[Safonov et. al. 2012; Rabitz et. al. 2005; Langbein et. al. 2015; Schirmer et. al. (2001, 2022); Floether et. al. 2012; Bukov et. al. 2018; Knill et. al. 2007; Goldschmidt et. al. 2022, …]

Big picture contributions

$$H(\mathbf{u}(t), t) = H_0 + H_c(\mathbf{u}(t), t)$$ + noise

1. Uncertain model: The controller learns the model *ab initio* by collecting {system, control, fidelity} data via RL during the control loop. (CDC 2021)

$$\frac{d}{dt}U(\mathbf{u}(t), t) = -\frac{i}{\hbar}H(\mathbf{u}(t), t)U(\mathbf{u}(t), t), \quad U(t = t_0) = \mathrm{I}$$

+ environmental decay

2. Resource intensive: Relax the full dynamics' learnability constraint by encoding a partial learnable model in the controller. (PRR 2023)

$$\mathcal{F}(U_{\text{target}}, U(t_0, t_1, \mathbf{u}(t))) = \frac{1}{2^{2n}}\left|\mathrm{Tr}\,\left[U^\dagger_{\text{target}}U(t_0, t_1, \mathbf{u}(t))\right]\right|^2$$

Is its optimal value stable under perturbations?

3. Fidelity missing stability: Generalise to RIM defined within a local region. Certify robustness. (PRA 2023)
4. Classical robust control connection?: Related to the differential-sensitivity. (CDC+CSL 2023)

2. Robustness certification

Optimality is not enough

1. Fidelity alone is not enough. Need ‘stable fidelity’ controllers.
2. Randomise the fidelity functional using a noise scale
3. Compute probabilistic distance w.r.t. Ideal value

Prop. (informal): **The robustness infidelity measure (RIM_p) is**

$$\mathbb{E}_{\mathbf{P}(\mathcal{F}=f)}\left[(1-f)^p\right]^{1/p}$$

4. p=1 produces an interesting interpretation of the mean infidelity
5. Correlates with other robust fidelity ‘metrics’

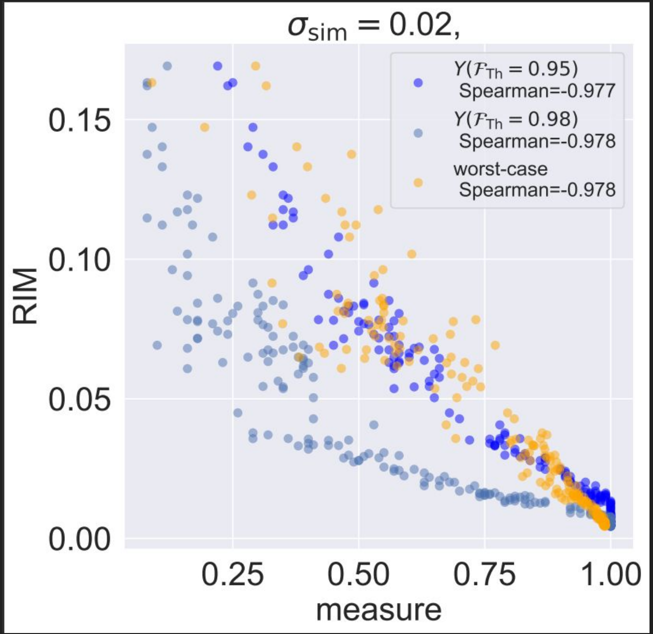

Compare controllers individually or as families

Ranked controllers by fidelity clearly over-optimise w.r.t. RIM

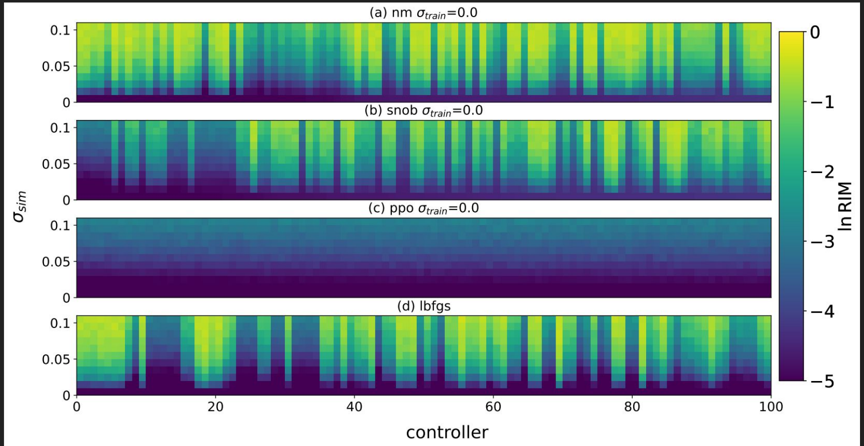

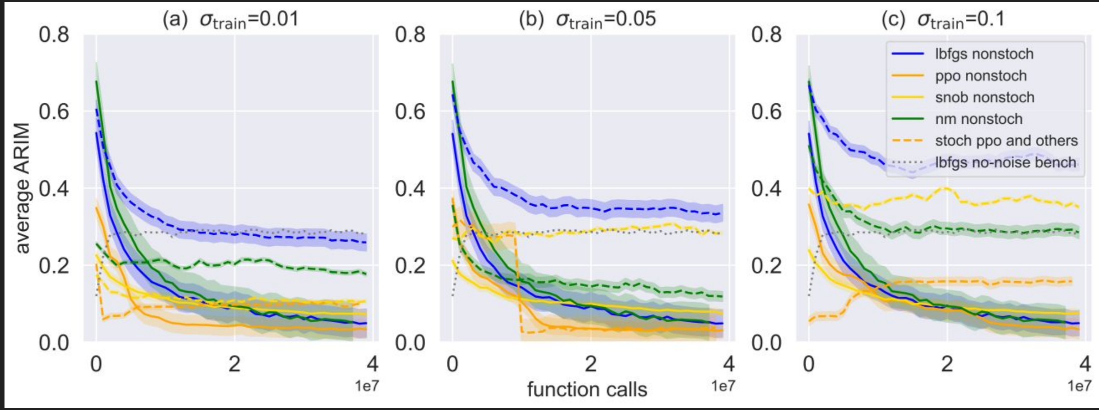

Can compare and optimise the RIM over control algo.

Model-free RL has better resource complexity under noise

[Jonckheere et. al. 2018, O’Neil et. al. 2023, Schirmer et. al. 2022]

Relation to classical robust control

Prop. (informal): **The expected differential sensitivity is the differential sensitivity of RIM$_1$**

$$\mathbb{E}_{\mathbf{P}(\mathbf{S}=S)}[\zeta(S_\mu, T)] = \left.\frac{\partial \mathrm{RIM}_1(\sigma)}{\partial \sigma}\right|_{\sigma=0}$$

where $$\zeta(S_\mu, T) = \left.\frac{\partial(1-\mathcal{F})(T; S_\mu, \sigma)}{\partial \sigma}\right|_{\sigma=0}$$

The classical robust control literature is very mature. Can be explored more.

3. Optimisation of quantum dynamics

Reinforcement learning without a model

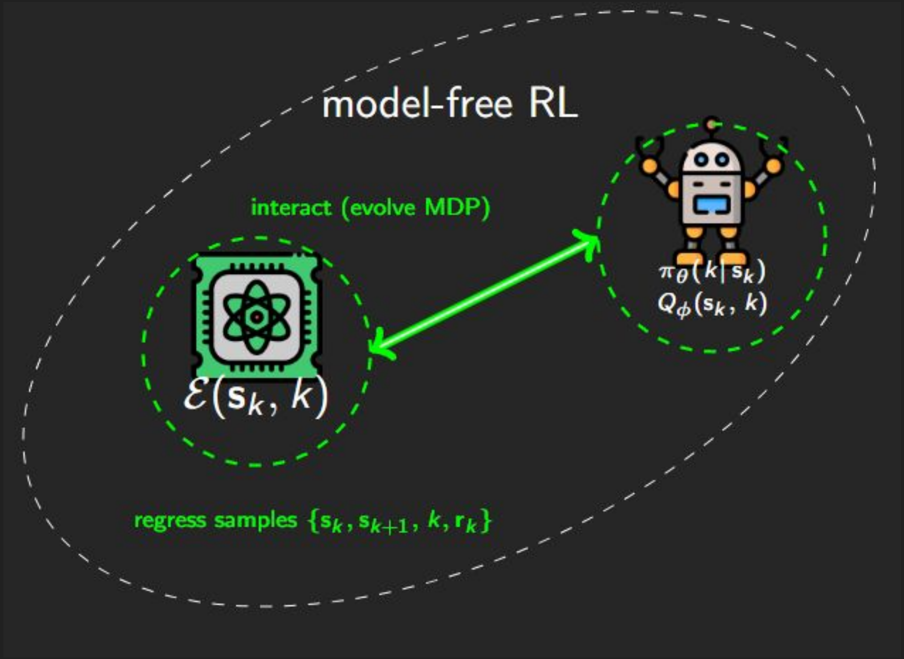

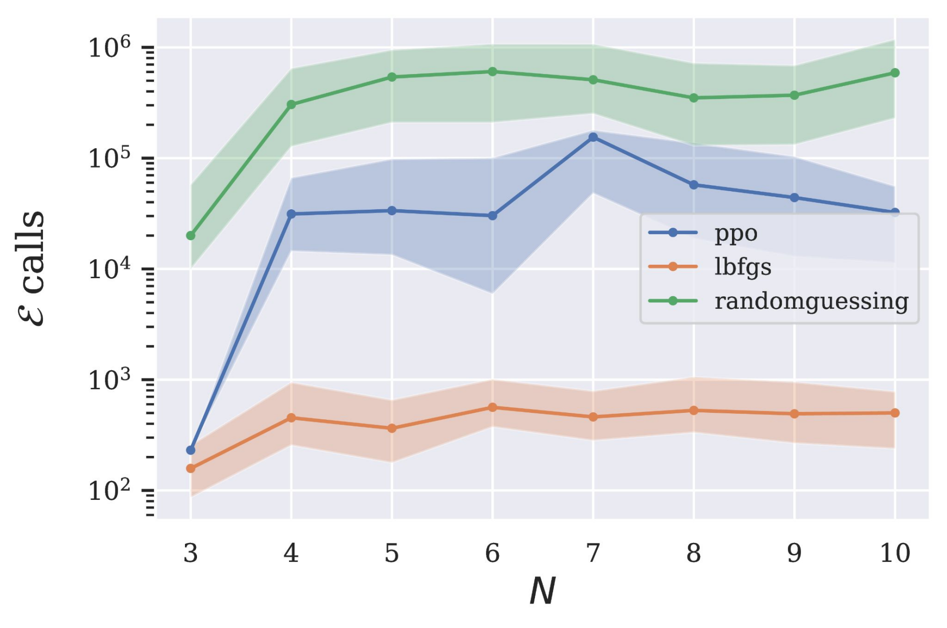

resource intensity hierarchy:
random > RL > model+gradient based

Learning for control: improving RL resource complexity

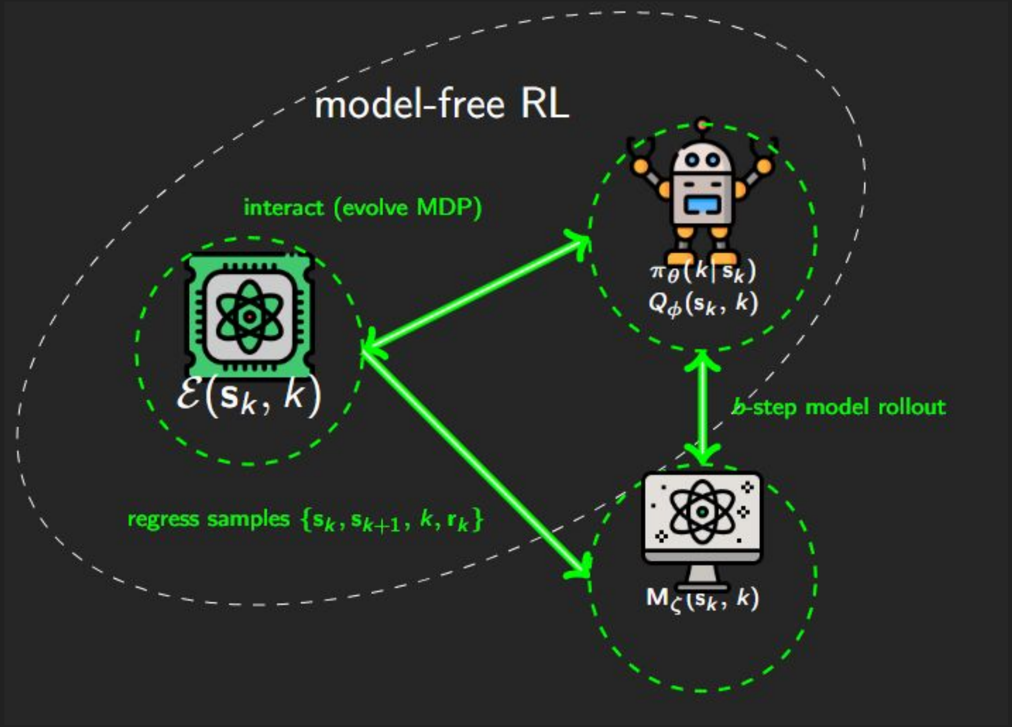

[Haarnoja et. al. 2018; Janner et. al. 2019; Kidger et. al. 2022; R.T.Q. Chen et. al. (2018), …,]

Model based RL (LH-MBSAC) learns a partial model directly using MDP data and terminates early

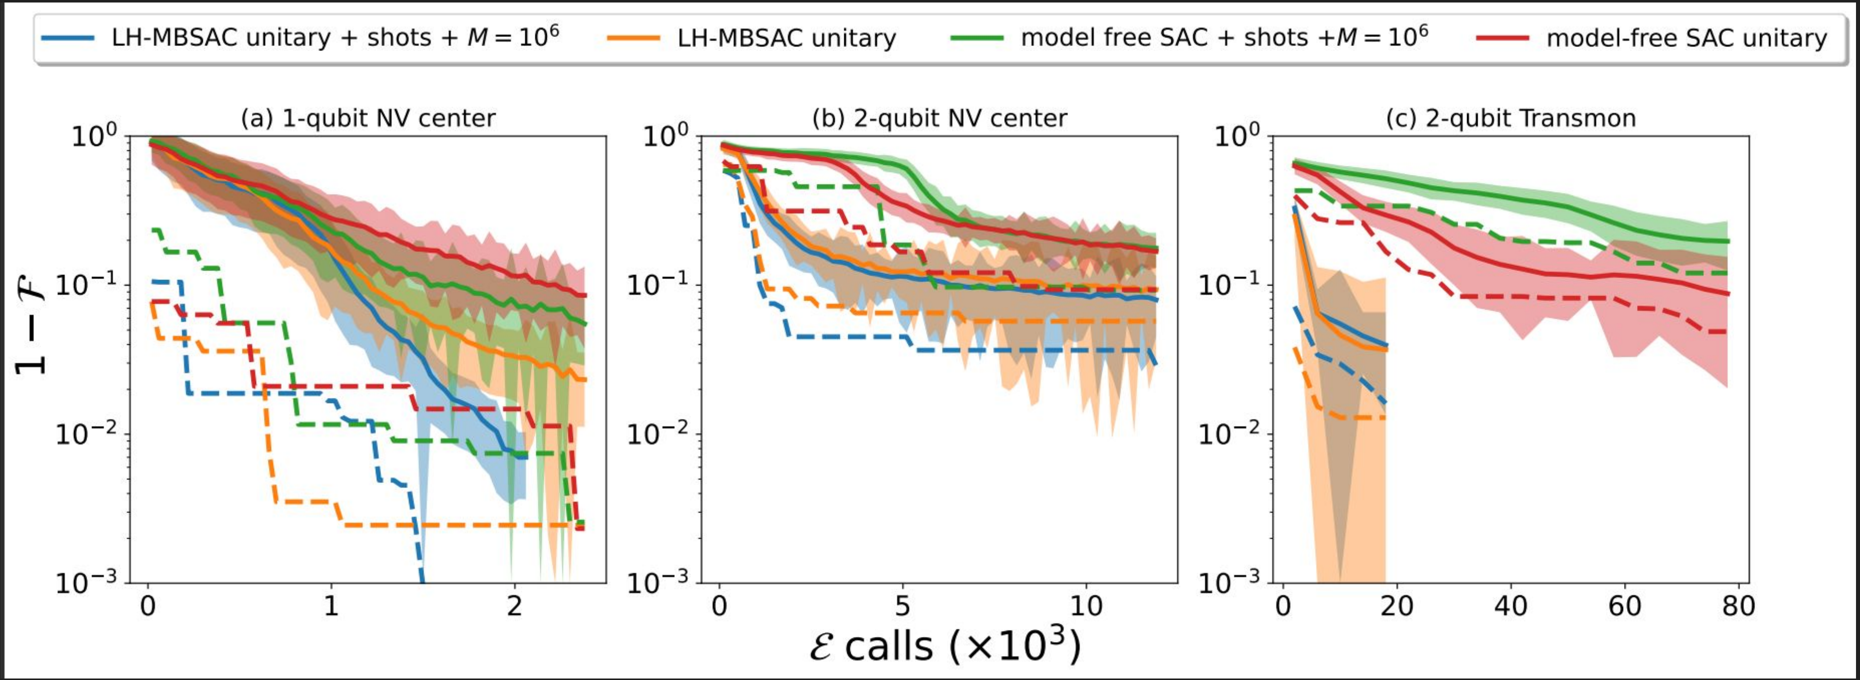

The learned model is approximately correct but useful

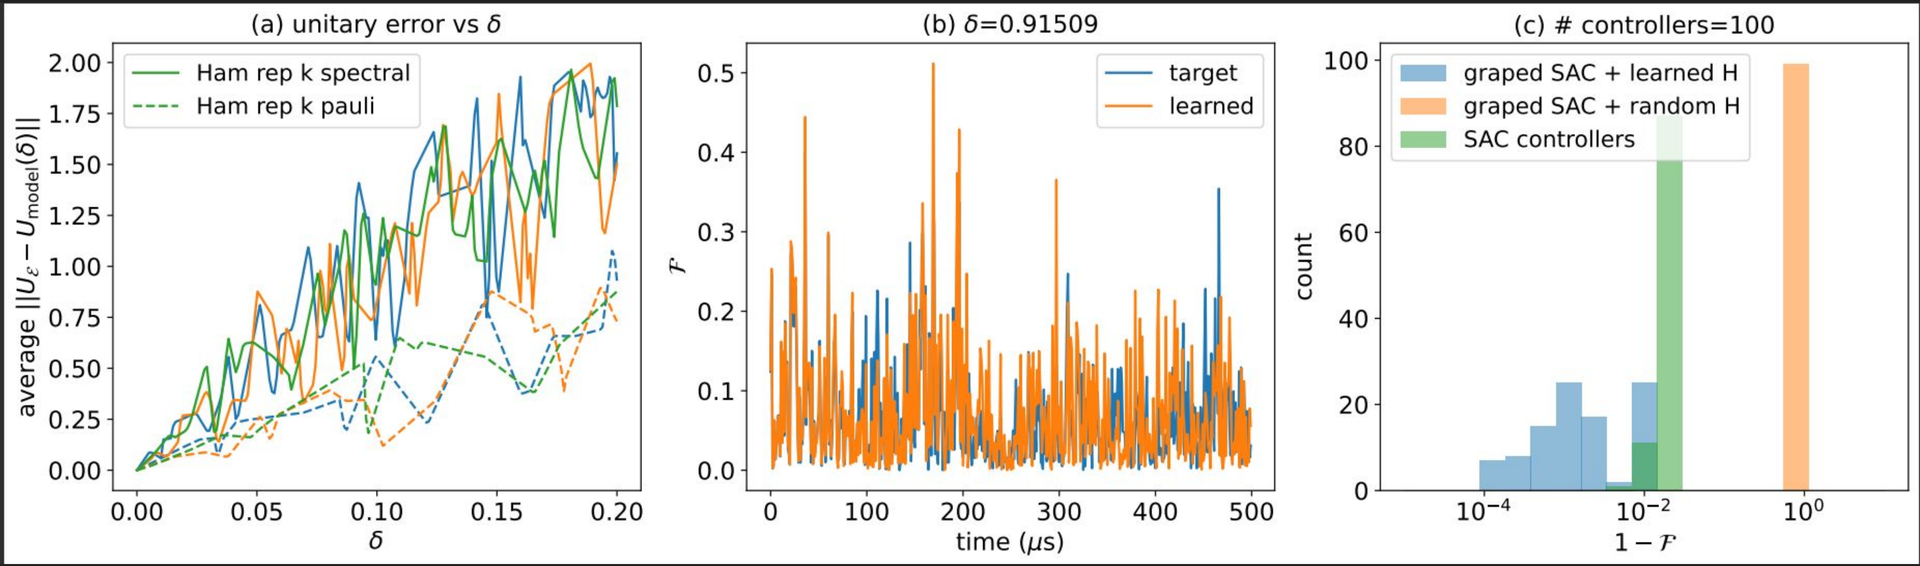

Learned-model based optimisation preserves controller robustness

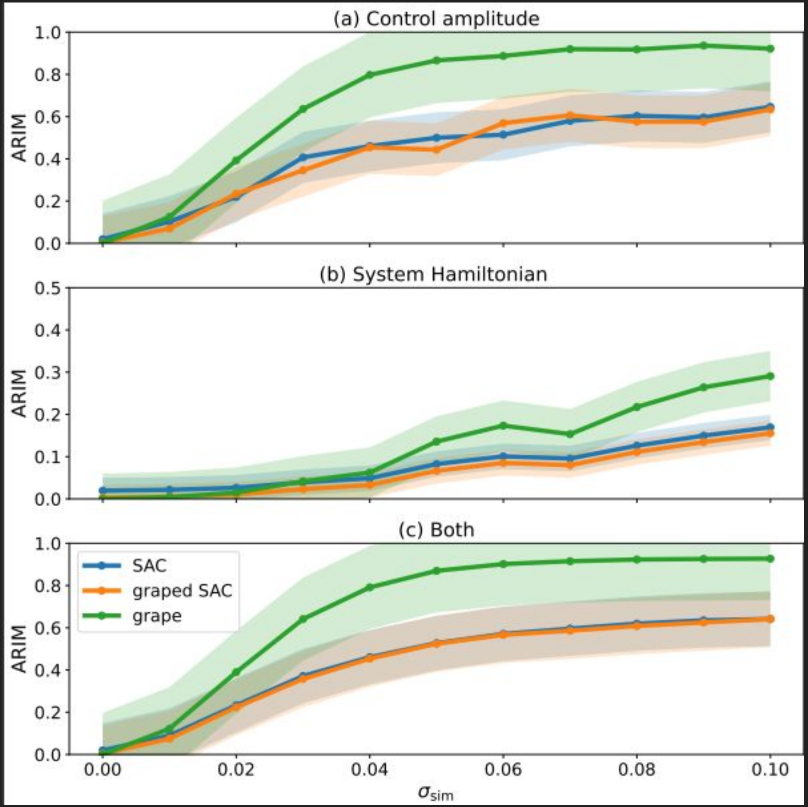

Acknowledgements






Fin

Backup: Structure within the RIM orders

Proposition 5.3. *The following bounds hold:*

$$RIM_{p'} \leqslant n^{\left(\frac{1}{p}-\frac{1}{p'}\right)} RIM_p, \tag{5.12a}$$

$$RIM_p \leqslant RIM_{p'} \tag{5.12b}$$

for $p < p'$, where n is the number of samples used to estimate the RIM.

- Convergence between different orders p, p’ at small magnitudes RIM → 0
- Sublinear growth in the regime: RIM >> 0
- Wasserstein distance is a probability structure preserving geodesic that facilitates convergence of different orders
- Large p measure is more sensitive to outliers
  - Hurts if measure is a comparator (outliers ‘blur’ the aggregated picture)
  - Helps if measure is a target (optimisation actively selects against outliers)