Quantum-enhanced adaptive agents with efficient long-term memories

Thomas J. Elliott

University of Manchester

November 24, 2023

The University of Manchester

Contents

Modelling a Stochastic Process

Suppose we have a stochastic process $\ldots X_{-3} X_{-2} X_{-1} X_0 X_1 X_2 X_3 \ldots$

Task: (Statistically) replicate the future behaviour of the process $P(\overrightarrow{X} | \overleftarrow{x})$

Modelling a Stochastic Process

Suppose we have a stochastic process $\ldots X_{-3}X_{-2}X_{-1}X_0X_1X_2X_3\ldots$

Task: (Statistically) replicate the future behaviour of the process $P(\overrightarrow{X}|\overleftarrow{x})$

Storing the entire past is infeasible...

...we must extract the useful information

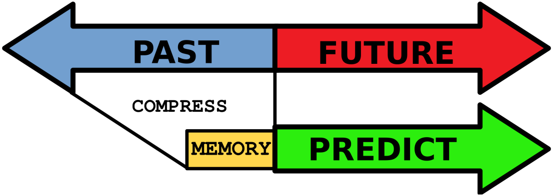

Modelling a Stochastic Process

Suppose we have a stochastic process $\ldots X_{-3}X_{-2}X_{-1}X_0X_1X_2X_3\ldots$

Task: (Statistically) replicate the future behaviour of the process $P(\overrightarrow{X}|\overleftarrow{x})$

Storing the entire past is infeasible...
...we must extract the useful information

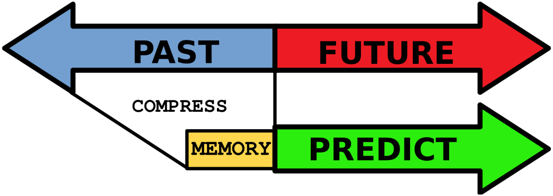

Compression by encoding function $f : \overleftarrow{\mathcal{X}} \to \sigma_{\mathcal{M}}$

Update function to produce outputs $\Lambda : \sigma_{\mathcal{M}} \to \sigma_{\mathcal{M}} \times \mathcal{X}$

Modelling a Stochastic Process

Memory cost: $C_f = -\mathrm{Tr}(\rho \log_2[\rho])$ (Note: $\rho = \sum_m P(m)\sigma_m$)

C. R. Shalizi and J. P. Crutchfield, J. Stat. Phys. **104** 817 (2001)

Modelling a Stochastic Process

Memory cost: $C_f = -\mathrm{Tr}(\rho \log_2[\rho])$ (Note: $\rho = \sum_m P(m)\sigma_m$)

Causal state encoding function

$$f_\varepsilon(\overleftarrow{x}) = f_\varepsilon(\overleftarrow{x}') \Leftrightarrow P(\overrightarrow{X}|\overleftarrow{X} = \overleftarrow{x}) = P(\overrightarrow{X}|\overleftarrow{X} = \overleftarrow{x}')$$

Provably memory-minimal classical encoding

C. R. Shalizi and J. P. Crutchfield, J. Stat. Phys. **104** 817 (2001)

Modelling a Stochastic Process

Memory cost: $C_f = -\mathrm{Tr}(\rho \log_2[\rho])$ (Note: $\rho = \sum_m P(m)\sigma_m$)

Causal state encoding function

$$f_\varepsilon(\overleftarrow{x}) = f_\varepsilon(\overleftarrow{x}') \Leftrightarrow P(\overrightarrow{X}|\overleftarrow{X} = \overleftarrow{x}) = P(\overrightarrow{X}|\overleftarrow{X} = \overleftarrow{x}')$$

Provably memory-minimal classical encoding

These minimal classical models still store redundant information

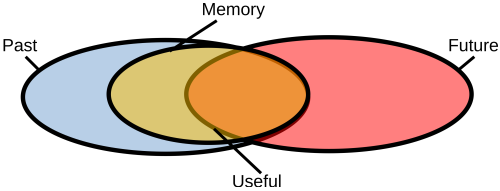

C. R. Shalizi and J. P. Crutchfield, J. Stat. Phys. **104** 817 (2001)

Modelling a Stochastic Process

Memory cost: $C_f = -\mathrm{Tr}(\rho \log_2[\rho])$ (Note: $\rho = \sum_m P(m)\sigma_m$)

Causal state encoding function

$$f_\varepsilon(\overleftarrow{x}) = f_\varepsilon(\overleftarrow{x}') \Leftrightarrow P(\overrightarrow{X}|\overleftarrow{X} = \overleftarrow{x}) = P(\overrightarrow{X}|\overleftarrow{X} = \overleftarrow{x}')$$

Provably memory-minimal classical encoding

These minimal classical models still store redundant information

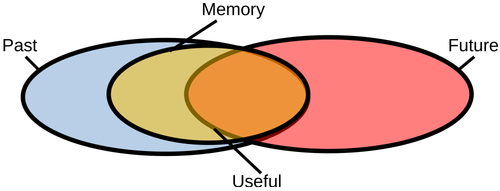

Fully distinguishable memory states give rise to partially distinguishable futures

$$f_\varepsilon : \overleftarrow{\mathcal{X}} \rightarrow \{|j\rangle\}$$

C. R. Shalizi and J. P. Crutchfield, J. Stat. Phys. **104** 817 (2001)

Quantum Compression Advantage

Quantum encodings can mitigate some of this redundancy

$$f_q : \overleftarrow{\mathcal{X}} \rightarrow \{|\sigma_j\rangle\}$$

M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Nat. Comm. **3** 762 (2012)
T. J. Elliott and M. Gu, npj Quantum Information **4** 18 (2018)
K.-D. Wu et al., Nature Communications **14** 2624 (2023)

Quantum Compression Advantage

Quantum encodings can mitigate some of this redundancy

$$f_q : \overleftarrow{\mathcal{X}} \rightarrow \{|\sigma_j\rangle\}$$

Whenever the ε-machine stores redundant information...

...a quantum model can do better!

$$C_\mu > I(\overleftarrow{X};\overrightarrow{X}) \quad \Leftrightarrow \quad C_q < C_\mu$$

M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Nat. Comm. **3** 762 (2012)

T. J. Elliott and M. Gu, npj Quantum Information **4** 18 (2018)

K.-D. Wu et al., Nature Communications **14** 2624 (2023)

Quantum Compression Advantage

Quantum encodings can mitigate some of this redundancy

$$f_q : \overleftarrow{\mathcal{X}} \rightarrow \{|\sigma_j\rangle\}$$

Whenever the ε-machine stores redundant information...

...a quantum model can do better!

$$C_\mu > I(\overleftarrow{X}; \overrightarrow{X}) \quad \Leftrightarrow \quad C_q < C_\mu$$

Scaling advantage

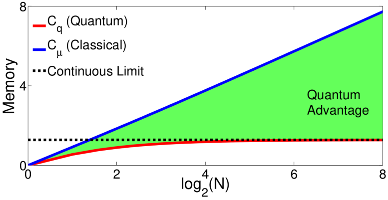

Experimentally implemented

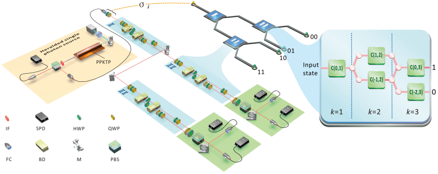

M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Nat. Comm. **3** 762 (2012)
T. J. Elliott and M. Gu, npj Quantum Information **4** 18 (2018)
K.-D. Wu et al., Nature Communications **14** 2624 (2023)

Adaptive Agents

What about systems that modify behaviour in response to environmental input?

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Adaptive Agents

What about systems that modify behaviour in response to environmental input?

For example, AI, self-driving cars, chatbots, and trading algorithms

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Adaptive Agents

What about systems that modify behaviour in response to environmental input?

For example, AI, self-driving cars, chatbots, and trading algorithms

These are **adaptive agents**

Process replaced by a *strategy*

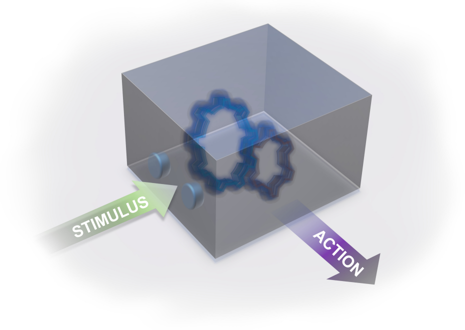

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Adaptive Agents

What about systems that modify behaviour in response to environmental input?

For example, AI, self-driving cars, chatbots, and trading algorithms

These are **adaptive agents**

Process replaced by a *strategy*

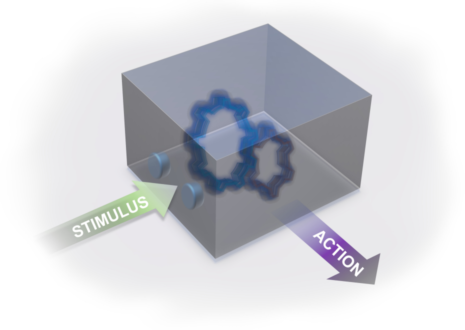

There is a trade-off between strategy complexity and memory cost

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Adaptive Agents

What about systems that modify behaviour in response to environmental input?

For example, AI, self-driving cars, chatbots, and trading algorithms

These are **adaptive agents**

Process replaced by a *strategy*

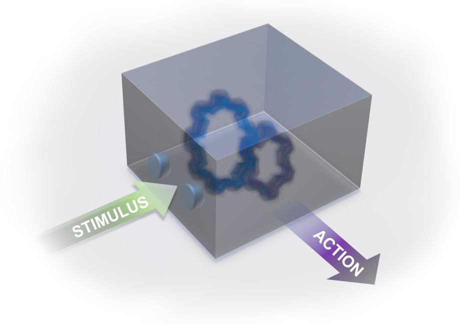

There is a trade-off between strategy complexity and memory cost

Can quantum technologies provide a competitive edge?

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ Stimuli (inputs) the agent can recognise

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ Stimuli (inputs) the agent can recognise
- $\mathcal{Y}$ Actions (outputs) the agent can perform

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ Stimuli (inputs) the agent can recognise
- $\mathcal{Y}$ Actions (outputs) the agent can perform
- $\{\sigma_m\}$ The agent's memory states

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ — Stimuli (inputs) the agent can recognise
- $\mathcal{Y}$ — Actions (outputs) the agent can perform
- $\{\sigma_m\}$ — The agent's memory states
- $f : \overleftarrow{\mathcal{X}} \times \overleftarrow{\mathcal{Y}} \to \{\sigma_m\}$ — The agent's memory encoding function

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ — Stimuli (inputs) the agent can recognise
- $\mathcal{Y}$ — Actions (outputs) the agent can perform
- $\{\sigma_m\}$ — The agent's memory states
- $f : \overleftarrow{\mathcal{X}} \times \overleftarrow{\mathcal{Y}} \to \{\sigma_m\}$ — The agent's memory encoding function
- $\Lambda : \mathcal{X} \times \{\sigma_m\} \to \mathcal{Y} \times \{\sigma_m\}$ — The agent's policy function

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ — Stimuli (inputs) the agent can recognise
- $\mathcal{Y}$ — Actions (outputs) the agent can perform
- $\{\sigma_m\}$ — The agent's memory states
- $f : \overleftarrow{\mathcal{X}} \times \overleftarrow{\mathcal{Y}} \to \{\sigma_m\}$ — The agent's memory encoding function
- $\Lambda : \mathcal{X} \times \{\sigma_m\} \to \mathcal{Y} \times \{\sigma_m\}$ — The agent's policy function

These enable the agent to implement a strategy $P(Y|\overleftarrow{X}, \overleftarrow{Y}, X)$

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ — Stimuli (inputs) the agent can recognise
- $\mathcal{Y}$ — Actions (outputs) the agent can perform
- $\{\sigma_m\}$ — The agent's memory states
- $f : \overleftarrow{\mathcal{X}} \times \overleftarrow{\mathcal{Y}} \rightarrow \{\sigma_m\}$ — The agent's memory encoding function
- $\Lambda : \mathcal{X} \times \{\sigma_m\} \rightarrow \mathcal{Y} \times \{\sigma_m\}$ — The agent's policy function

These enable the agent to implement a strategy $P(Y|\overleftarrow{X}, \overleftarrow{Y}, X)$

Minimal classical agents:

$$f_\varepsilon(\overleftarrow{z}) = f_\varepsilon(\overleftarrow{z}') \Leftrightarrow P(\overrightarrow{Y}|\overleftarrow{z}, \overrightarrow{x}) = P(\overrightarrow{Y}|\overleftarrow{z}', \overrightarrow{x}) \forall \overrightarrow{x}$$

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Anatomy of an Adaptive Agent

An adaptive agent is defined by:

- $\mathcal{X}$ — Stimuli (inputs) the agent can recognise
- $\mathcal{Y}$ — Actions (outputs) the agent can perform
- $\{\sigma_m\}$ — The agent's memory states
- $f : \overleftarrow{\mathcal{X}} \times \overleftarrow{\mathcal{Y}} \rightarrow \{\sigma_m\}$ — The agent's memory encoding function
- $\Lambda : \mathcal{X} \times \{\sigma_m\} \rightarrow \mathcal{Y} \times \{\sigma_m\}$ — The agent's policy function

These enable the agent to implement a strategy $P(Y|\overleftarrow{X}, \overleftarrow{Y}, X)$

Minimal classical agents:

$$f_\varepsilon(\overleftarrow{z}) = f_\varepsilon(\overleftarrow{z}') \Leftrightarrow P(\overrightarrow{Y}|\overleftarrow{z}, \overrightarrow{x}) = P(\overrightarrow{Y}|\overleftarrow{z}', \overrightarrow{x}) \forall \overrightarrow{x}$$

Quantum agents: $\{\sigma_m\}$ are quantum states, Λ a quantum channel

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Quantum-Enhanced Adaptive Agents

Provably memory-minimal form of a quantum adaptive agent:

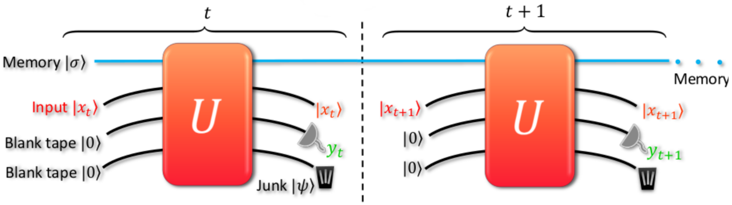

$$U|\sigma_s\rangle|x\rangle|0\rangle|0\rangle = \sum_y \sqrt{P(y|x,s)}|\sigma_{\lambda(z,s)}\rangle|x\rangle|y\rangle|\psi(z,s)\rangle$$

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Quantum-Enhanced Adaptive Agents

Provably memory-minimal form of a quantum adaptive agent:

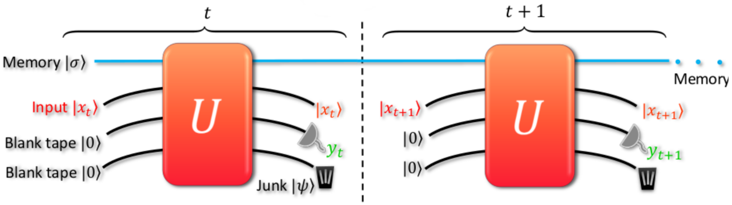

$$U|\sigma_s\rangle|x\rangle|0\rangle|0\rangle = \sum_y \sqrt{P(y|x,s)}|\sigma_{\lambda(z,s)}\rangle|x\rangle|y\rangle|\psi(z,s)\rangle$$

Memory states pure and one-to-one with causal states

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Quantum-Enhanced Adaptive Agents

Provably memory-minimal form of a quantum adaptive agent:

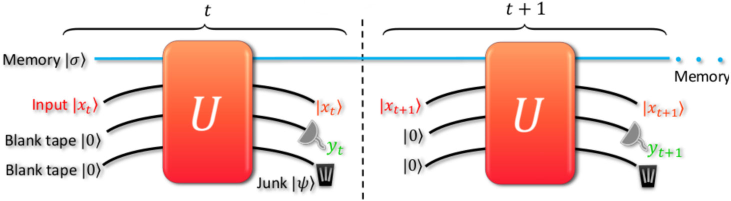

$$U|\sigma_s\rangle|x\rangle|0\rangle|0\rangle = \sum_y \sqrt{P(y|x,s)}|\sigma_{\lambda(z,s)}\rangle|x\rangle|y\rangle|\psi(z,s)\rangle$$

Memory states pure and one-to-one with causal states

Classical inputs and outputs

Quantum processing only *within* agent

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Quantum-Enhanced Adaptive Agents

Algorithm Systematic quantum agent encoding

Inputs: Causal states $\mathcal{S}$, transition probabilities $P(Y|X,S)$, and update rule $\lambda(z,s)$
Outputs: Quantum memory states $\{|\sigma_s\rangle\}$, evolution operator U

1: Construct the set of multivariate polynomial equations

$$c^x_{ss'} = \sum_y \sqrt{P(y|x,s)P(y|x,s')} \prod_{x'} c^{x'}_{\lambda(z,s)\lambda(z,s')} \tag{1}$$

defined $\forall s, s' \in \mathcal{S}, x \in \mathcal{X}$ and solve to obtain $\{c^x_{ss'}\}$

2: Use a reverse Gram-Schmidt procedure to construct quantum memory states $\{|\sigma_s\rangle\}$ from overlaps $c_{ss'} = \prod_x c^x_{ss'}$, and junk states $\{|\psi(z,s)\rangle\}$ from overlaps $d^z_{ss'} = \prod_{x' \neq x} c^{x'}_{ss'}$

3: Construct the columns of U explicitly defined

4: Use a Gram-Schmidt procedure to fill the remaining columns of U, ensuring orthogonality with existing columns

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Quantum-Enhanced Adaptive Agents

Algorithm Systematic quantum agent encoding

Inputs: Causal states $\mathcal{S}$, transition probabilities $P(Y|X,S)$, and update rule $\lambda(z,s)$
Outputs: Quantum memory states $\{|\sigma_s\rangle\}$, evolution operator U

1: Construct the set of multivariate polynomial equations

$$c^x_{ss'} = \sum_y \sqrt{P(y|x,s)P(y|x,s')} \prod_{x'} c^{x'}_{\lambda(z,s)\lambda(z,s')} \tag{1}$$

defined $\forall s, s' \in \mathcal{S}, x \in \mathcal{X}$ and solve to obtain $\{c^x_{ss'}\}$

2: Use a reverse Gram-Schmidt procedure to construct quantum memory states $\{|\sigma_s\rangle\}$ from overlaps $c_{ss'} = \prod_x c^x_{ss'}$, and junk states $\{|\psi(z,s)\rangle\}$ from overlaps $d^z_{ss'} = \prod_{x' \neq x} c^{x'}_{ss'}$

3: Construct the columns of U explicitly defined

4: Use a Gram-Schmidt procedure to fill the remaining columns of U, ensuring orthogonality with existing columns

Quantum advantage unless for all possible pairs of states $\exists \overrightarrow{x}$ that leads to perfectly distinguishable future action sequences

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Example: Resettable stochastic clocks

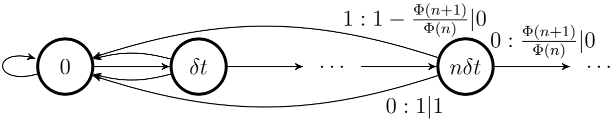

$\Phi(t)$: Probability of 'survival' to time t without ticking *under natural evolution*
δt: Resolution of the coarse-grained timesteps

Inputs
$x = 0$ *Evolve naturally*
$x = 1$ *Reset*

Outputs
$x = 0$ *No Tick*
$x = 1$ *Tick*

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Example: Resettable stochastic clocks

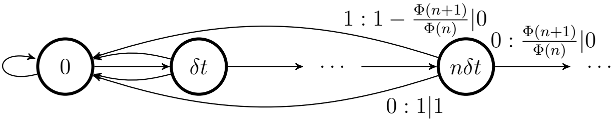

$\Phi(t)$: Probability of 'survival' to time t without ticking *under natural evolution*
δt: Resolution of the coarse-grained timesteps

Inputs		**Outputs**	
$x = 0$	*Evolve naturally*	$x = 0$	*No Tick*
$x = 1$	*Reset*	$x = 1$	*Tick*

We expect memory cost to increase with $1/\delta t$

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Example: Resettable stochastic clocks

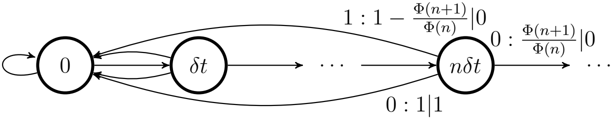

$\Phi(t)$: Probability of 'survival' to time t without ticking *under natural evolution*
δt: Resolution of the coarse-grained timesteps

Inputs
$x = 0$ *Evolve naturally*
$x = 1$ *Reset*

Outputs
$x = 0$ *No Tick*
$x = 1$ *Tick*

We expect memory cost to increase with $1/\delta t$

True for classical agents...
... but not for quantum!

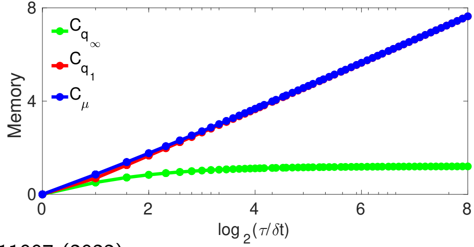

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Scaling advantages can be found in more general settings

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Scaling advantages can be found in more general settings

Consider *n*-bit discretisation of (finite) continuous parameter τ into $\delta\tau^{(n)}$

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Scaling advantages can be found in more general settings

Consider n-bit discretisation of (finite) continuous parameter τ into $\delta\tau^{(n)}$

- **Distributional convergence:** Memory state steady-state probability (densities) converge exponentially with increasing precision

$$|P^{(n)}(\tau^{(n)})/\delta\tau^{(n)} - P^{(n-1)}(\tau^{(n)})/\delta\tau^{(n-1)}| < K\delta\tau^{(n)}$$

Refined states approximately share equal weighting

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Scaling advantages can be found in more general settings

Consider n-bit discretisation of (finite) continuous parameter τ into $\delta\tau^{(n)}$

- **Distributional convergence:** Memory state steady-state probability (densities) converge exponentially with increasing precision

$$|P^{(n)}(\tau^{(n)})/\delta\tau^{(n)} - P^{(n-1)}(\tau^{(n)})/\delta\tau^{(n-1)}| < K\delta\tau^{(n)}$$

  Refined states approximately share equal weighting

- **Memory-overlap convergence:** Memory state overlaps converge exponentially with increasing precision

$$|c^{(n)}_{\tau^{(n)}\tau'^{(n)}} - c^{(n-1)}_{\tau^{(n)}\tau'^{(n)}}| < K\delta\tau^{(n)}$$

  where $c_{jk} := \langle\sigma_j|\sigma_k\rangle$
  Refined states have very similar distributions

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Scaling advantages can be found in more general settings

Consider n-bit discretisation of (finite) continuous parameter τ into $\delta\tau^{(n)}$

- **Distributional convergence:** Memory state steady-state probability (densities) converge exponentially with increasing precision

$$|P^{(n)}(\tau^{(n)})/\delta\tau^{(n)} - P^{(n-1)}(\tau^{(n)})/\delta\tau^{(n-1)}| < K\delta\tau^{(n)}$$

  Refined states approximately share equal weighting

- **Memory-overlap convergence:** Memory state overlaps converge exponentially with increasing precision

$$|c^{(n)}_{\tau^{(n)}\tau'^{(n)}} - c^{(n-1)}_{\tau^{(n)}\tau'^{(n)}}| < K\delta\tau^{(n)}$$

  where $c_{jk} := \langle\sigma_j|\sigma_k\rangle$
  Refined states have very similar distributions

If these conditions are satisfied, C_q remains bounded as $n \to \infty$

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Scalable Quantum Advantage

Scaling advantages can be found in more general settings

Consider n-bit discretisation of (finite) continuous parameter τ into $\delta\tau^{(n)}$

- **Distributional convergence:** Memory state steady-state probability (densities) converge exponentially with increasing precision

$$|P^{(n)}(\tau^{(n)})/\delta\tau^{(n)} - P^{(n-1)}(\tau^{(n)})/\delta\tau^{(n-1)}| < K\delta\tau^{(n)}$$

Refined states approximately share equal weighting

- **Memory-overlap convergence:** Memory state overlaps converge exponentially with increasing precision

$$|c^{(n)}_{\tau^{(n)}\tau'^{(n)}} - c^{(n-1)}_{\tau^{(n)}\tau'^{(n)}}| < K\delta\tau^{(n)}$$

where $c_{jk} := \langle\sigma_j|\sigma_k\rangle$

Refined states have very similar distributions

If these conditions are satisfied, C_q remains bounded as $n \to \infty$

Stochastic clocks: Satisfied if $\Phi(t)$ infinitely differentiable and $\sim e^{-\gamma t}$ at long times

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)

Thermodynamical Advantage for Quantum Agents

Quantum memory advantages correspond to reduced thermodynamical footprint

Thermodynamical Advantage for Quantum Agents

Quantum memory advantages correspond to reduced thermodynamical footprint

Agent must expend energy to clean up junk

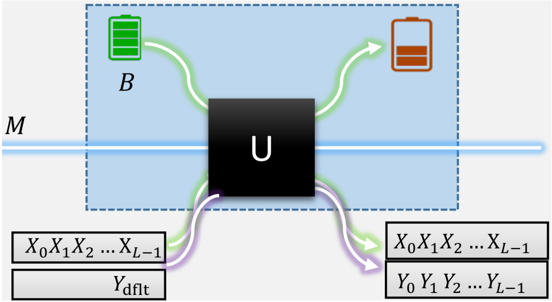

Thermodynamical Advantage for Quantum Agents

Quantum memory advantages correspond to reduced thermodynamical footprint

Agent must expend energy to clean up junk

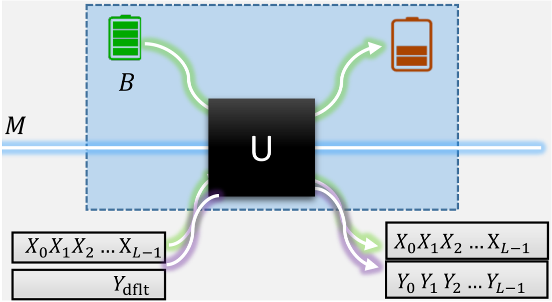

For agents in parallel, work cost (per agent) to process L inputs at once given by

$$\frac{\Delta W}{k_{\mathrm{b}} T \ln 2} = S(Y_{\mathrm{in}0:L}) + S(X_{0:L}) + S(M_0) - S(X_{0:L}, Y_{0:L}, M_L)$$

Thermodynamical Advantage for Quantum Agents

Assume i.i.d. inputs, steady-state memory; work rate (per timestep cost) becomes

$$\frac{w_L}{k_\mathrm{b} T \ln 2} = S(Y_\mathrm{in}) + \frac{1}{L}[I(X_{0:L}, Y_{0:L}; M_L) - S(Y_{0:L}|X_{0:L})]$$

Thermodynamical Advantage for Quantum Agents

Assume i.i.d. inputs, steady-state memory; work rate (per timestep cost) becomes

$$\frac{w_L}{k_{\mathrm{b}} T \ln 2} = S(Y_{\mathrm{in}}) + \frac{1}{L}[I(X_{0:L}, Y_{0:L}; M_L) - S(Y_{0:L}|X_{0:L})]$$

Quantum advantage given by

$$w_L^c - w_L^q = \frac{k_{\mathrm{b}} T \ln 2}{L}[I(X_{0:L}, Y_{0:L}; M_L^c) - I(X_{0:L}, Y_{0:L}; M_L^q)]$$

Thermodynamical Advantage for Quantum Agents

Assume i.i.d. inputs, steady-state memory; work rate (per timestep cost) becomes

$$\frac{w_L}{k_{\mathrm{b}} T \ln 2} = S(Y_{\mathrm{in}}) + \frac{1}{L}[I(X_{0:L}, Y_{0:L}; M_L) - S(Y_{0:L}|X_{0:L})]$$

Quantum advantage given by

$$w_L^c - w_L^q = \frac{k_{\mathrm{b}} T \ln 2}{L}[I(X_{0:L}, Y_{0:L}; M_L^c) - I(X_{0:L}, Y_{0:L}; M_L^q)]$$

By data processing inequality, $w_L^c - w_L^q > 0$ if $C_q < C_\mu$

Thermodynamical Advantage for Quantum Agents

Assume i.i.d. inputs, steady-state memory; work rate (per timestep cost) becomes

$$\frac{w_L}{k_{\mathrm{b}} T \ln 2} = S(Y_{\mathrm{in}}) + \frac{1}{L}[I(X_{0:L}, Y_{0:L}; M_L) - S(Y_{0:L}|X_{0:L})]$$

Quantum advantage given by

$$w_L^c - w_L^q = \frac{k_{\mathrm{b}} T \ln 2}{L}[I(X_{0:L}, Y_{0:L}; M_L^c) - I(X_{0:L}, Y_{0:L}; M_L^q)]$$

By data processing inequality, $w_L^c - w_L^q > 0$ if $C_q < C_\mu$

For offline processing ($L \to \infty$)

$$w^c - w^q \propto C_\mu - C_q$$

Thus, analogous scalable quantum advantage in thermal efficiency

Summary & Future

Summary:

- Quantum computers can simulate complex processes with less memory
- Framework can be extended to quantum-enhanced adaptive agents
- Such quantum advantages can scale with increasing complexity
- Superior thermal efficiency of quantum agents

Summary & Future

Summary:

- Quantum computers can simulate complex processes with less memory
- Framework can be extended to quantum-enhanced adaptive agents
- Such quantum advantages can scale with increasing complexity
- Superior thermal efficiency of quantum agents

Future:

- Quantum advantage in memory dimension
- Applications of memory-efficient quantum agents
- Incorporation with quantum ‘speed-ups’ for agents

Thanks for listening!

Mile Gu

Andrew Garner

Jayne Thompson

MANCHESTER 1824 The University of Manchester

FQXi Foundational Questions Institute

Imperial College London

Nanyang Technological University Singapore

T. J. Elliott et al., Phys. Rev. X **12** 011007 (2022)