Confidence intervals for the mixing time of a reversible Markov chain from a single sample path

Daniel Hsu† Aryeh Kontorovich♯ Csaba Szepesvári⋆

†Columbia University, ♯Ben-Gurion University, ⋆University of Alberta

ITA 2016

Problem

- Irreducible, aperiodic, time-homogeneous Markov chain

$$X_1 \rightarrow X_2 \rightarrow X_3 \rightarrow \cdots$$

Problem

- Irreducible, aperiodic, time-homogeneous Markov chain

$$X_1 \;\to\; X_2 \;\to\; X_3 \;\to\; \cdots$$

- There is a unique stationary distribution π with

$$\lim_{t\to\infty} \mathcal{L}(X_t \mid X_1 = x) \;=\; \pi\,, \quad \text{for all } x \in \mathcal{X}\,.$$

Problem

- Irreducible, aperiodic, time-homogeneous Markov chain

$$X_1 \;\rightarrow\; X_2 \;\rightarrow\; X_3 \;\rightarrow\; \cdots$$

- There is a unique stationary distribution π with

$$\lim_{t\to\infty} \mathcal{L}(X_t \mid X_1 = x) \;=\; \pi\,, \quad \text{for all } x \in \mathcal{X}\,.$$

- The mixing time t_{mix} is the earliest time t with

$$\sup_{x\in\mathcal{X}} \|\mathcal{L}(X_t \mid X_1 = x) - \pi\|_{\text{tv}} \;\leq\; 1/4\,.$$

Problem

- Irreducible, aperiodic, time-homogeneous Markov chain

$$X_1 \;\to\; X_2 \;\to\; X_3 \;\to\; \cdots$$

- There is a unique stationary distribution π with

$$\lim_{t\to\infty} \mathcal{L}(X_t \mid X_1 = x) \;=\; \pi\,, \quad \text{for all } x \in \mathcal{X}\,.$$

- The mixing time t_{mix} is the earliest time t with

$$\sup_{x\in\mathcal{X}} \|\mathcal{L}(X_t \mid X_1 = x) - \pi\|_{\mathrm{tv}} \;\leq\; 1/4\,.$$

Problem:

Determine (confidently) if $t \geq t_{\mathrm{mix}}$ after seeing $X_1, X_2, \ldots, X_t$.

Problem

- Irreducible, aperiodic, time-homogeneous Markov chain

$$X_1 \to X_2 \to X_3 \to \cdots$$

- There is a unique stationary distribution π with

$$\lim_{t\to\infty} \mathcal{L}(X_t \mid X_1 = x) = \pi\,, \quad \text{for all } x \in \mathcal{X}\,.$$

- The mixing time t_{mix} is the earliest time t with

$$\sup_{x\in\mathcal{X}} \|\mathcal{L}(X_t \mid X_1 = x) - \pi\|_{\mathrm{tv}} \le 1/4\,.$$

Problem:

Given $\delta \in (0,1)$ and $X_{1:t}$, determine non-trivial $I_t \subseteq [0,\infty]$ with

$$\mathbb{P}(t_{\mathrm{mix}} \in I_t) \ge 1 - \delta\,.$$

Some motivation from machine learning and statistics

Chernoff bounds for Markov chains $X_1 \to X_2 \to \cdots$:
for suitably well-behaved $f\colon \mathcal{X} \to \mathbb{R}$, with probability at least $1 - \delta$,

$$\left| \frac{1}{t} \sum_{i=1}^{t} f(X_i) - \mathbb{E}_\pi f \right| \;\le\; \underbrace{\tilde{O}\left(\sqrt{\frac{t_{\mathrm{mix}} \log(1/\delta)}{t}} \right)}_{\text{deviation bound}} .$$

Bound depends on t_{mix}, which may be unknown *a priori*.

Some motivation from machine learning and statistics

Chernoff bounds for Markov chains $X_1 \to X_2 \to \cdots$:
for suitably well-behaved $f\colon \mathcal{X} \to \mathbb{R}$, with probability at least $1-\delta$,

$$\left| \frac{1}{t} \sum_{i=1}^{t} f(X_i) - \mathbb{E}_\pi f \right| \;\leq\; \underbrace{\tilde{O}\left(\sqrt{\frac{t_{\mathrm{mix}} \log(1/\delta)}{t}} \right)}_{\text{deviation bound}} .$$

Bound depends on t_{mix}, which may be unknown *a priori*.

Examples:

Bayesian inference Posterior means & variances via MCMC

Reinforcement learning Mean action rewards in an MDP

Supervised learning Error rates of hypotheses from non-iid data

Some motivation from machine learning and statistics

Chernoff bounds for Markov chains $X_1 \to X_2 \to \cdots$:
for suitably well-behaved $f \colon \mathcal{X} \to \mathbb{R}$, with probability at least $1 - \delta$,

$$\left| \frac{1}{t} \sum_{i=1}^{t} f(X_i) - \mathbb{E}_\pi f \right| \;\leq\; \underbrace{\tilde{O}\left(\sqrt{\frac{t_{\mathrm{mix}} \log(1/\delta)}{t}} \right)}_{\text{deviation bound}} .$$

Bound depends on t_{mix}, which may be unknown *a priori*.

Examples:

Bayesian inference Posterior means & variances via MCMC

Reinforcement learning Mean action rewards in an MDP

Supervised learning Error rates of hypotheses from non-iid data

Need *observable* deviation bounds.

Observable deviation bounds from mixing time bounds?

Suppose an estimator $\hat{t}_{\mathrm{mix}} = \hat{t}_{\mathrm{mix}}(X_{1:t})$ of t_{mix} satisfies:

$$\mathbb{P}(t_{\mathrm{mix}} \leq \hat{t}_{\mathrm{mix}} + \varepsilon_t) \ \geq \ 1 - \delta \, .$$

Observable deviation bounds from mixing time bounds?

Suppose an estimator $\hat{t}_{\mathsf{mix}} = \hat{t}_{\mathsf{mix}}(X_{1:t})$ of t_{mix} satisfies:

$$\mathbb{P}(t_{\mathsf{mix}} \le \hat{t}_{\mathsf{mix}} + \varepsilon_t) \;\ge\; 1 - \delta\,.$$

Then with probability at least $1 - 2\delta$,

$$\left| \frac{1}{t} \sum_{i=1}^{t} f(X_i) - \mathbb{E}_\pi f \right| \;\le\; \tilde{O}\left(\sqrt{\frac{(\hat{t}_{\mathsf{mix}} + \varepsilon_t) \log(1/\delta)}{t}} \right).$$

Observable deviation bounds from mixing time bounds?

Suppose an estimator $\hat{t}_{\mathrm{mix}} = \hat{t}_{\mathrm{mix}}(X_{1:t})$ of t_{mix} satisfies:

$$\mathbb{P}(t_{\mathrm{mix}} \leq \hat{t}_{\mathrm{mix}} + \varepsilon_t) \;\geq\; 1 - \delta\,.$$

Then with probability at least $1 - 2\delta$,

$$\left| \frac{1}{t} \sum_{i=1}^{t} f(X_i) - \mathbb{E}_\pi f \right| \;\leq\; \tilde{O}\left(\sqrt{\frac{(\hat{t}_{\mathrm{mix}} + \varepsilon_t) \log(1/\delta)}{t}} \right).$$

But $\hat{t}_{\mathrm{mix}}$ is computed from $X_{1:t}$, so ε_t may also depend on t_{mix}.

Observable deviation bounds from mixing time bounds?

Suppose an estimator $\hat{t}_{\text{mix}} = \hat{t}_{\text{mix}}(X_{1:t})$ of t_{mix} satisfies:

$$\mathbb{P}(t_{\text{mix}} \leq \hat{t}_{\text{mix}} + \varepsilon_t) \;\geq\; 1 - \delta \,.$$

Then with probability at least $1 - 2\delta$,

$$\left| \frac{1}{t} \sum_{i=1}^{t} f(X_i) - \mathbb{E}_\pi f \right| \;\leq\; \tilde{O}\left(\sqrt{\frac{(\hat{t}_{\text{mix}} + \varepsilon_t) \log(1/\delta)}{t}} \right).$$

But $\hat{t}_{\text{mix}}$ is computed from $X_{1:t}$, so ε_t may also depend on t_{mix}.

Deviation bounds for point estimators are insufficient.
Need *(observable) confidence intervals* for t_{mix}.

What we do

What we do

1. Shift focus to relaxation time t_{relax} to enable spectral methods.

What we do

1. Shift focus to relaxation time t_{relax} to enable spectral methods.
2. **Lower/upper bounds** on sample path length for point estimation of t_{relax}.

What we do

1. Shift focus to relaxation time t_{relax} to enable spectral methods.
2. **Lower/upper bounds** on sample path length for point estimation of t_{relax}.
3. **New algorithm** for constructing confidence intervals for t_{relax}.

Relaxation time

- Let P be the transition operator of the Markov chain, and let $\lambda_\star$ be its second-largest eigenvalue modulus (i.e., largest eigenvalue modulus other than 1).

Relaxation time

- Let P be the transition operator of the Markov chain, and let $\lambda_\star$ be its second-largest eigenvalue modulus (i.e., largest eigenvalue modulus other than 1).
- Spectral gap: $\gamma_\star := 1 - \lambda_\star$.
  Relaxation time: $t_{\mathsf{relax}} := 1/\gamma_\star$.

$$(t_{\mathsf{relax}} - 1)\ln 2 \;\le\; t_{\mathsf{mix}} \;\le\; t_{\mathsf{relax}} \ln \frac{4}{\pi_\star}$$

for $\pi_\star := \min_{x \in \mathcal{X}} \pi(x)$.

Relaxation time

- Let P be the transition operator of the Markov chain, and let $\lambda_\star$ be its second-largest eigenvalue modulus (i.e., largest eigenvalue modulus other than 1).
- Spectral gap: $\gamma_\star := 1 - \lambda_\star$.
  Relaxation time: $t_{\text{relax}} := 1/\gamma_\star$.

$$(t_{\text{relax}} - 1) \ln 2 \;\leq\; t_{\text{mix}} \;\leq\; t_{\text{relax}} \ln \frac{4}{\pi_\star}$$

  for $\pi_\star := \min_{x \in \mathcal{X}} \pi(x)$.

  Assumptions on P ensure $\gamma_\star, \pi_\star \in (0,1)$.

Relaxation time

- Let P be the transition operator of the Markov chain, and let $\lambda_\star$ be its second-largest eigenvalue modulus (i.e., largest eigenvalue modulus other than 1).
- Spectral gap: $\gamma_\star := 1 - \lambda_\star$.
  Relaxation time: $t_{\mathsf{relax}} := 1/\gamma_\star$.

$$(t_{\mathsf{relax}} - 1)\ln 2 \;\le\; t_{\mathsf{mix}} \;\le\; t_{\mathsf{relax}} \ln \frac{4}{\pi_\star}$$

  for $\pi_\star := \min_{x \in \mathcal{X}} \pi(x)$.

  Assumptions on P ensure $\gamma_\star, \pi_\star \in (0,1)$.

Spectral approach: construct CI's for $\gamma_\star$ and $\pi_\star$.

Our results (point estimation)

We restrict to reversible Markov chains on finite state spaces.

Let d be the (known *a priori*) cardinality of the state space $\mathcal{X}$.

Our results (point estimation)

We restrict to reversible Markov chains on finite state spaces.
Let d be the (known *a priori*) cardinality of the state space $\mathcal{X}$.

1. **Lower bound**:
   To estimate $\gamma_\star$ within a constant multiplicative factor, every algorithm needs (w.p. $1/4$) sample path of length

$$\geq \Omega\left(\frac{d \log d}{\gamma_\star} + \frac{1}{\pi_\star}\right).$$

Our results (point estimation)

We restrict to reversible Markov chains on finite state spaces.
Let d be the (known *a priori*) cardinality of the state space $\mathcal{X}$.

1. **Lower bound**:
   To estimate $\gamma_\star$ within a constant multiplicative factor, every algorithm needs (w.p. 1/4) sample path of length

$$\geq \Omega\left(\frac{d \log d}{\gamma_\star} + \frac{1}{\pi_\star}\right).$$

2. **Upper bound**:
   Simple algorithm estimates $\gamma_\star$ and $\pi_\star$ within a constant multiplicative factor (w.h.p.) with sample path of length

$$\widetilde{O}\left(\frac{\log d}{\pi_\star \gamma_\star^3}\right) \quad (\text{for } \gamma_\star), \qquad \widetilde{O}\left(\frac{\log d}{\pi_\star \gamma_\star}\right) \quad (\text{for } \pi_\star).$$

Our results (point estimation)

We restrict to reversible Markov chains on finite state spaces.
Let d be the (known *a priori*) cardinality of the state space $\mathcal{X}$.

1. **Lower bound**:
   To estimate $\gamma_\star$ within a constant multiplicative factor, every algorithm needs (w.p. 1/4) sample path of length
   $$\geq \ \Omega\left(\frac{d \log d}{\gamma_\star} + \frac{1}{\pi_\star}\right).$$
2. **Upper bound**:
   Simple algorithm estimates $\gamma_\star$ and $\pi_\star$ within a constant multiplicative factor (w.h.p.) with sample path of length
   $$\widetilde{O}\left(\frac{\log d}{\pi_\star \gamma_\star^3}\right) \quad (\text{for } \gamma_\star), \qquad \widetilde{O}\left(\frac{\log d}{\pi_\star \gamma_\star}\right) \quad (\text{for } \pi_\star).$$

But point estimator $\not\Rightarrow$ confidence interval.

Our results (confidence intervals)

3. **New algorithm**: Given $\delta \in (0,1)$ and $X_{1:t}$ as input, constructs intervals $I_t^{\gamma_\star}$ and $I_t^{\pi_\star}$ such that

$$\mathbb{P}\big(\gamma_\star \in I_t^{\gamma_\star}\big) \;\geq\; 1-\delta \quad \text{and} \quad \mathbb{P}\big(\pi_\star \in I_t^{\pi_\star}\big) \;\geq\; 1-\delta\,.$$

   Widths of intervals converge a.s. to zero at $\sqrt{\frac{\log\log t}{t}}$ rate.

Our results (confidence intervals)

3. **New algorithm**: Given $\delta \in (0,1)$ and $X_{1:t}$ as input, constructs intervals $I_t^{\gamma_\star}$ and $I_t^{\pi_\star}$ such that

$$\mathbb{P}\left(\gamma_\star \in I_t^{\gamma_\star}\right) \;\geq\; 1-\delta \quad \text{and} \quad \mathbb{P}\left(\pi_\star \in I_t^{\pi_\star}\right) \;\geq\; 1-\delta\,.$$

   Widths of intervals converge a.s. to zero at $\sqrt{\frac{\log\log t}{t}}$ rate.

4. **Hybrid approach**: Use new algorithm to turn error bounds for point estimators into observable CI's.
   (This improves asymptotic rate for $\pi_\star$ interval.)

Plug-in estimator

Plug-in estimator

- Reversibility grants the symmetry of

$$M \;:=\; \mathrm{diag}(\pi)P \;=\; \left\{\mathbb{P}_{X_1\sim\pi}(X_1 = x, X_2 = x')\right\}_{x,x'\in\mathcal{X}}$$

(doublet state probabilities in stationary chain).

Plug-in estimator

- Reversibility grants the symmetry of

$$M \;:=\; \mathrm{diag}(\pi)P \;=\; \left\{\mathbb{P}_{X_1\sim\pi}(X_1 = x, X_2 = x')\right\}_{x,x'\in\mathcal{X}}$$

(doublet state probabilities in stationary chain).

- Moreover, eigenvalues of

$$L \;:=\; \mathrm{diag}(\pi)^{-1/2} M\, \mathrm{diag}(\pi)^{-1/2}$$

are real, and satisfy

$$1 \;=\; \lambda_1 \;>\; \lambda_2 \;\geq\; \cdots \;\geq\; \lambda_d \;>\; -1\,,$$
$$\gamma_\star \;=\; 1 - \max\{\lambda_2, |\lambda_d|\}\,.$$

Plug-in estimator

- Reversibility grants the symmetry of

$$M \;:=\; \operatorname{diag}(\pi)P \;=\; \left\{\mathbb{P}_{X_1\sim\pi}(X_1 = x, X_2 = x')\right\}_{x,x'\in\mathcal{X}}$$

(doublet state probabilities in stationary chain).

- Moreover, eigenvalues of

$$L \;:=\; \operatorname{diag}(\pi)^{-1/2} M \operatorname{diag}(\pi)^{-1/2}$$

are real, and satisfy

$$1 \;=\; \lambda_1 \;>\; \lambda_2 \;\geq\; \cdots \;\geq\; \lambda_d \;>\; -1\,,$$
$$\gamma_\star \;=\; 1 - \max\{\lambda_2, |\lambda_d|\}\,.$$

- **Plug-in estimator**: estimate π and M from $X_{1:t}$ (using empirical frequencies), then plug-in to formula for $\gamma_\star$.

Chicken-and-egg problem

(Matrix) Chernoff bound (for Markov chains) gives error bounds for estimates of π and M (and ultimately of L and $\gamma_\star$): e.g., w.h.p.,

$$|\hat{\gamma}_\star - \gamma_\star| \;\leq\; \|\widehat{L} - L\| \;\leq\; O\left(\sqrt{\frac{\log(d)\log(t/\pi_\star)}{\gamma_\star \pi_\star t}}\right).$$

Chicken-and-egg problem

(Matrix) Chernoff bound (for Markov chains) gives error bounds for estimates of π and M (and ultimately of L and $\gamma_\star$): e.g., w.h.p.,

$$|\hat{\gamma}_\star - \gamma_\star| \;\leq\; \|\widehat{L} - L\| \;\leq\; O\left(\sqrt{\frac{\log(d)\log(t/\pi_\star)}{\gamma_\star \pi_\star t}}\right).$$

This has *inverse* dependence on $\gamma_\star$.

Chicken-and-egg problem

(Matrix) Chernoff bound (for Markov chains) gives error bounds for estimates of π and M (and ultimately of L and $\gamma_\star$): e.g., w.h.p.,

$$|\hat{\gamma}_\star - \gamma_\star| \;\leq\; \|\widehat{L} - L\| \;\leq\; O\left(\sqrt{\frac{\log(d)\log(t/\pi_\star)}{\gamma_\star \pi_\star t}}\right).$$

This has *inverse* dependence on $\gamma_\star$.

Can't "solve the bound" for $\gamma_\star$

(unlike "empirical Bernstein" inequalities).

Direct estimation of P

Alternative: directly estimate P from $X_{1:t}$.

Direct estimation of P

Alternative: directly estimate P from $X_{1:t}$.

- **Key advantage**: observable confidence intervals for P via "empirical Bernstein" inequality for martingales.

Direct estimation of P

Alternative: directly estimate P from $X_{1:t}$.

- **Key advantage**: observable confidence intervals for P via "empirical Bernstein" inequality for martingales.

Two problems:

Direct estimation of P

Alternative: directly estimate P from $X_{1:t}$.

- **Key advantage**: observable confidence intervals for P via "empirical Bernstein" inequality for martingales.

Two problems:

1. Without appealing to symmetry structure, can argue

$$\|\widehat{P} - P\| \ \le \ \varepsilon \quad \Longrightarrow \quad |\hat{\gamma}_\star - \gamma_\star| \ \le \ O(\varepsilon^{1/(2d)}) \,,$$

but this implies exponential slow-down in rate.

Direct estimation of P

Alternative: directly estimate P from $X_{1:t}$.

- **Key advantage**: observable confidence intervals for P via "empirical Bernstein" inequality for martingales.

Two problems:

1. Without appealing to symmetry structure, can argue

$$\|\widehat{P} - P\| \;\leq\; \varepsilon \quad \Longrightarrow \quad |\hat{\gamma}_\star - \gamma_\star| \;\leq\; O(\varepsilon^{1/(2d)}),$$

   but this implies exponential slow-down in rate.

2. Direct appeal to symmetry structure of

$$L \;=\; \mathrm{diag}(\pi)^{1/2} P \,\mathrm{diag}(\pi)^{-1/2}$$

   gives bounds that depend on π, which is unknown.

Direct estimation of P

Alternative: directly estimate P from $X_{1:t}$.

- **Key advantage**: observable confidence intervals for P via "empirical Bernstein" inequality for martingales.

Two problems:

1. Without appealing to symmetry structure, can argue

$$\|\widehat{P} - P\| \le \varepsilon \quad \Longrightarrow \quad |\hat{\gamma}_\star - \gamma_\star| \le O(\varepsilon^{1/(2d)}),$$

   but this implies exponential slow-down in rate.

2. Direct appeal to symmetry structure of

$$L = \operatorname{diag}(\pi)^{1/2} P \operatorname{diag}(\pi)^{-1/2}$$

   gives bounds that depend on π, which is unknown.

Our approach:
Directly estimate P, and *indirectly* estimate π via $\widehat{P}$.

Indirect estimation of π

1. We ensure that $\widehat{P}$ is transition operator for an ergodic chain (easy via Laplace smoothing).

Indirect estimation of π

1. We ensure that $\widehat{P}$ is transition operator for an ergodic chain (easy via Laplace smoothing).
2. **Key step**: estimate π via $\widehat{P}$ via group inverse $\widehat{A}^{\#}$ of $I - \widehat{P}$.

Indirect estimation of π

1. We ensure that $\widehat{P}$ is transition operator for an ergodic chain (easy via Laplace smoothing).

2. **Key step**: estimate π via $\widehat{P}$ via group inverse $\widehat{A}^{\#}$ of $I - \widehat{P}$.

   - $\widehat{A}^{\#}$ contains "virtually everything that one would want to know about the chain" [with transition operator $\widehat{P}$] (Meyer, 1975).

Indirect estimation of π

1. We ensure that $\widehat{P}$ is transition operator for an ergodic chain (easy via Laplace smoothing).
2. **Key step**: estimate π via $\widehat{P}$ via group inverse $\widehat{A}^{\#}$ of $I - \widehat{P}$.
   - $\widehat{A}^{\#}$ contains "virtually everything that one would want to know about the chain" [with transition operator $\widehat{P}$] (Meyer, 1975).
   - Reveals unique stationary distribution $\hat{\pi}$ w.r.t. $\widehat{P}$.
     This is our indirect estimate of π.

Indirect estimation of π

1. We ensure that $\widehat{P}$ is transition operator for an ergodic chain (easy via Laplace smoothing).
2. **Key step**: estimate π via $\widehat{P}$ via group inverse $\widehat{A}^{\#}$ of $I - \widehat{P}$.
   - $\widehat{A}^{\#}$ contains "virtually everything that one would want to know about the chain" [with transition operator $\widehat{P}$] (Meyer, 1975).
   - Reveals unique stationary distribution $\hat{\pi}$ w.r.t. $\widehat{P}$. This is our indirect estimate of π.
   - Tells us how to bound $\|\hat{\pi} - \pi\|_\infty$ in terms of $\|\widehat{P} - P\|$. Hence, from this, we construct a confidence interval for π.

Overall algorithm (outline)

1. Form empirical estimate and confidence intervals for P
   (exploit Markov property & "empirical Bernstein"-type bounds).
2. Form estimate and confidence intervals for π
   (via group inverse of $I - \widehat{P}$).
3. Form estimate and confidence interval for $\gamma_\star$
   (via confidence intervals for π and P, & eigenvalue perturbation theory).

Recap and future work

- We resolve "chicken-and-egg" problem of observable confidence intervals for mixing time from a single sample path.

Recap and future work

- We resolve "chicken-and-egg" problem of observable confidence intervals for mixing time from a single sample path.
- Strongly exploit Markov property and ergodicity in confidence intervals for P and π.

Recap and future work

- We resolve "chicken-and-egg" problem of observable confidence intervals for mixing time from a single sample path.
- Strongly exploit Markov property and ergodicity in confidence intervals for P and π.
- **Problem #1**: close gap between lower and upper bounds on sample path length (for point estimation).

Recap and future work

- We resolve "chicken-and-egg" problem of observable confidence intervals for mixing time from a single sample path.
- Strongly exploit Markov property and ergodicity in confidence intervals for P and π.
- **Problem #1**: close gap between lower and upper bounds on sample path length (for point estimation).
- **Problem #2**: overcome computational bottlenecks from matrix operations.

Recap and future work

- We resolve "chicken-and-egg" problem of observable confidence intervals for mixing time from a single sample path.
- Strongly exploit Markov property and ergodicity in confidence intervals for P and π.
- **Problem #1**: close gap between lower and upper bounds on sample path length (for point estimation).
- **Problem #2**: overcome computational bottlenecks from matrix operations.
- **Problem #3**: handle large/continuous state spaces under suitable assumptions.

Recap and future work

- We resolve "chicken-and-egg" problem of observable confidence intervals for mixing time from a single sample path.
- Strongly exploit Markov property and ergodicity in confidence intervals for P and π.
- **Problem #1**: close gap between lower and upper bounds on sample path length (for point estimation).
- **Problem #2**: overcome computational bottlenecks from matrix operations.
- **Problem #3**: handle large/continuous state spaces under suitable assumptions.

Thanks!