Trace Inequivalence for MDPs is in P

Qiyi Tang
Joint work with Stefan Kiefer

SAMSA, Warsaw

Labelled Markov chain (LMC)

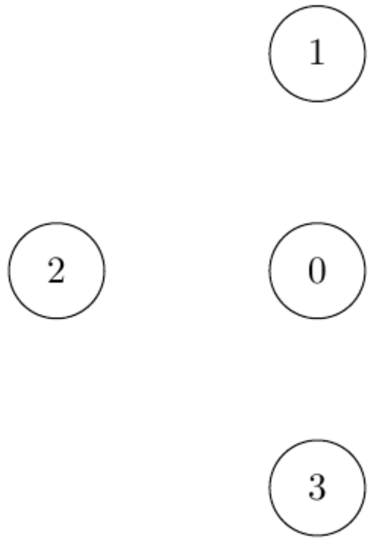

Labelled Markov chain (LMC)

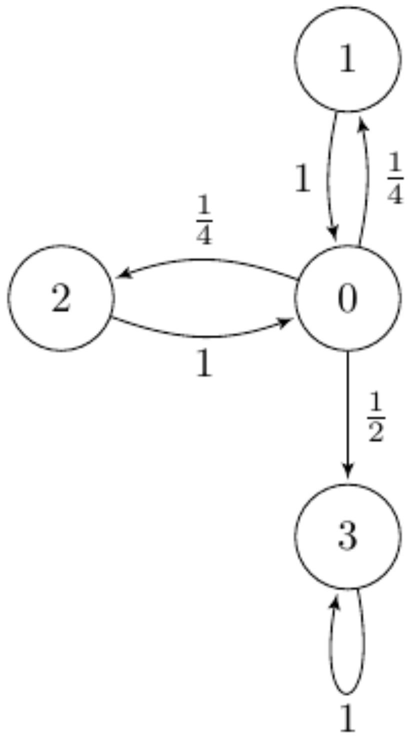

Labelled Markov chain (LMC)

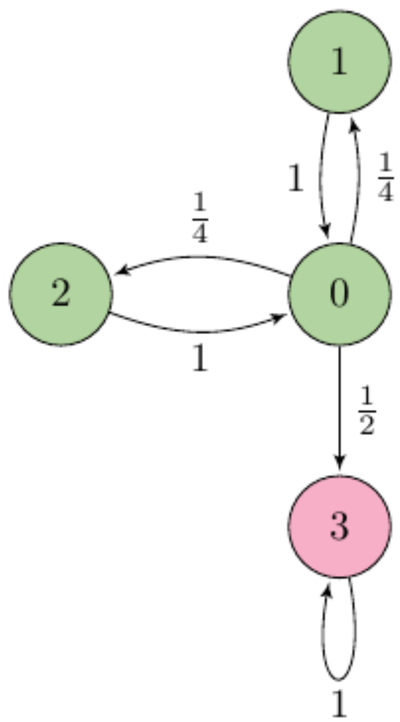

Trace probability

- Initial distribution $\mu = (1\ 0\ 0\ 0)$
- Transition matrices

$$M\ (\quad) = \begin{pmatrix} 0 & 1/4 & 1/4 & 1/2 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

$$M\ (\quad) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

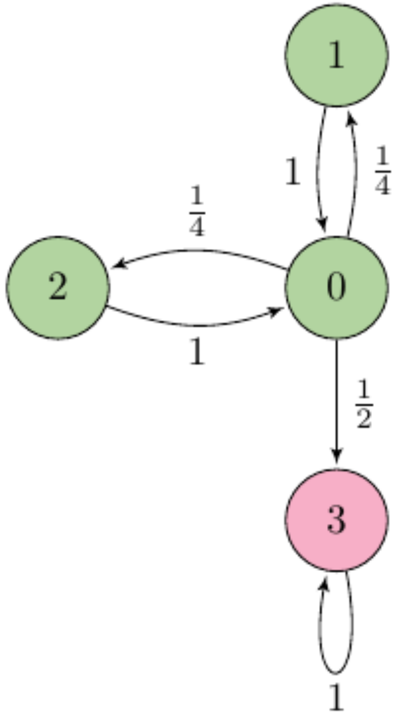

$M\ (\quad) + M\ (\quad)$ is stochastic

Trace probability

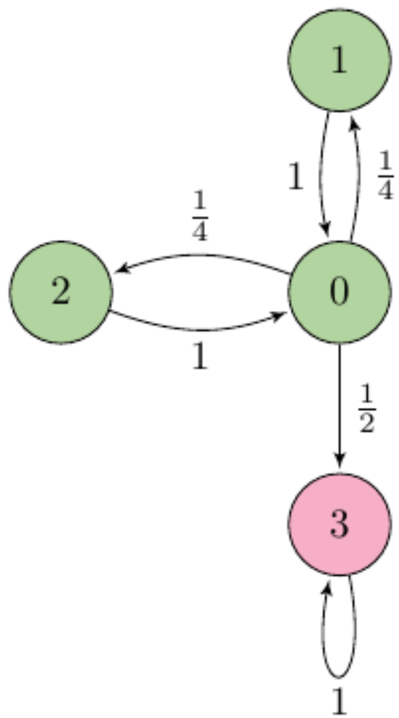

Q: What is the probability of the runs starting in state 0 that start with 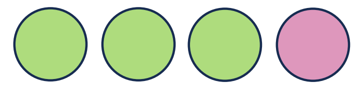 ?

Trace probability

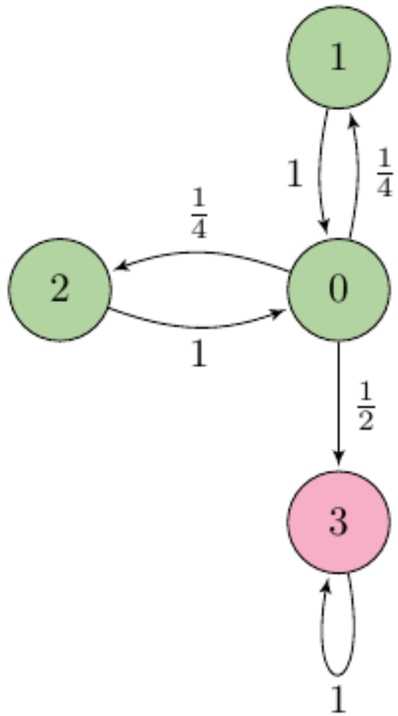

Q: What is the probability of the runs starting in state 0 that start with 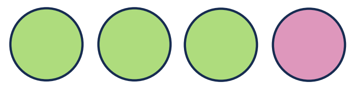 ?

A: 1/4 x 1/2 + 1/4 x 1/2

Trace probability

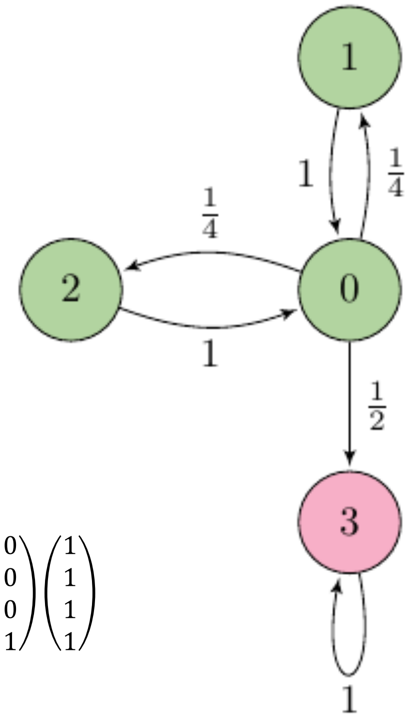

$\mu M(\bigcirc) M(\bigcirc)\ M(\bigcirc)\ M(\bigcirc)\vec{\mathbf{1}}$

$= \mu M(\bigcirc\bigcirc\bigcirc\bigcirc)\vec{\mathbf{1}}$

$= (1\ 0\ 0\ 0)\begin{pmatrix}0 & 1/4 & 1/4 & 1/2\\ 1 & 0 & 0 & 0\\ 1 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\end{pmatrix}\begin{pmatrix}0 & 1/4 & 1/4 & 1/2\\ 1 & 0 & 0 & 0\\ 1 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\end{pmatrix}\begin{pmatrix}0 & 1/4 & 1/4 & 1/2\\ 1 & 0 & 0 & 0\\ 1 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\end{pmatrix}\begin{pmatrix}0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 1\end{pmatrix}\begin{pmatrix}1\\ 1\\ 1\\ 1\end{pmatrix}$

=1/4

Q: What is the probability of the runs starting in state 0 that start with $\bigcirc\ \bigcirc\ \bigcirc\ \bigcirc$?

A: $1/4$ x $1/2 + 1/4$ x $1/2$

Trace (in)equivalence for LMCs

Two initial distributions μ_1 and μ_2 are trace equivalent

iff

the probabilities of generating all labelled sequences $w \in L^*$ are the same.

$$V = < M(w)\vec{\mathbf{1}} : w \in L^* >$$

Trace equivalence: $\mu_1 \equiv \mu_2$ iff $\mu_1 \vec{v} = \mu_2 \vec{v}$ for all $\vec{v} \in V$

Trace inequivalence: $\mu_1 \not\equiv \mu_2$ iff $\mu_1 \vec{v} \neq \mu_2 \vec{v}$ for some $\vec{v} \in V$

Known to have polynomial algorithms [Schützenberger, 1961]

Trace inequivalence problem for MDPs

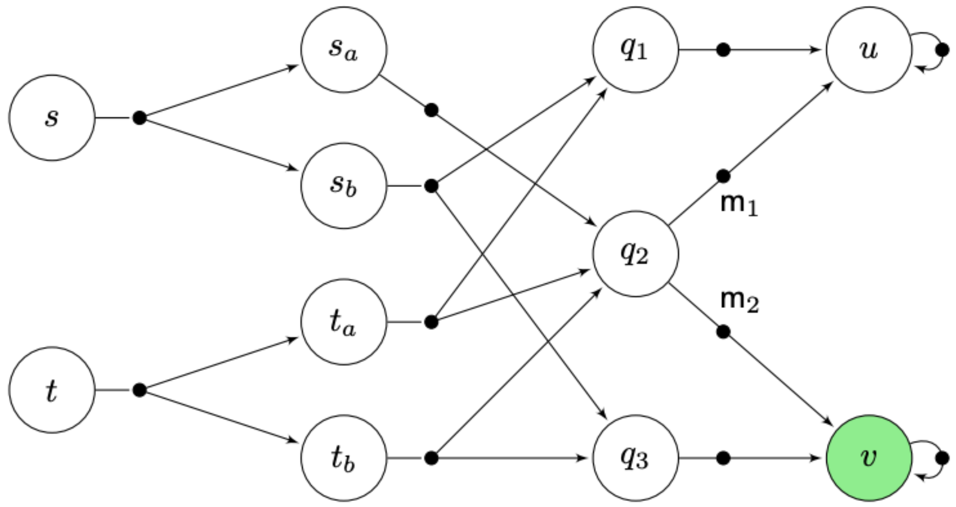

Given two initial distributions μ_1 and μ_2, is there a **strategy** α such that μ_1 and μ_2 are **not** trace equivalent, i.e., $\mu_1 \not\equiv \mu_2$, in the induced LMC.

- Memoryless deterministic (MD) strategies
- Memoryless strategies
- General strategies

Trace inequivalence problem for MDPs

$V = < M_{\alpha}(w)\vec{\mathbf{1}}: \alpha\ is\ a\ memoryless\ strategy, w \in L^{*} >$

$\mu_1 \not\equiv \mu_2$ iff $\mu_1 \vec{v} \neq \mu_2 \vec{v}$ for some $\vec{v} \in V$

Trace inequivalence problem for MDPs

$V_1 = < M_{\alpha_1}(w_1) M_{\alpha_2}(w_2) \ldots M_{\alpha_m}(w_m)\vec{\mathbf{1}} : \alpha_i \text{ is a memoryless strategy}, \ w = w_1 w_2 \ldots w_m \in L^* >$

$V_2 = < M_\alpha(w)\vec{\mathbf{1}} : \alpha \text{ is a memoryless strategy}, w \in L^* >$

$\mu_1 \not\equiv \mu_2$ iff $\mu_1 \vec{v} \neq \mu_2 \vec{v}$ for some $\vec{v} \in V_2$

$V_3 = < M_\alpha(w)\vec{\mathbf{1}} : \alpha \text{ is an MD strategy}, w \in L^* >$

$$V_1 \supseteq V_2 \supseteq V_3$$

Trace inequivalence problem for MDPs

$V_1 = < M_{\alpha_1}(w_1) M_{\alpha_2}(w_2) \ldots M_{\alpha_m}(w_m)\vec{\mathbf{1}}: \alpha_i \textit{ is a memoryless strategy}, \; w = w_1 w_2 \ldots w_m \in L^* >$

$V_2 = < M_\alpha(w)\vec{\mathbf{1}}: \alpha \textit{ is a memoryless strategy}, w \in L^* >$

$\mu_1 \not\equiv \mu_2$ iff $\mu_1 \vec{v} \neq \mu_2 \vec{v}$ for some $\vec{v} \in V_2$

$V_3 = < M_\alpha(w)\vec{\mathbf{1}}: \alpha \textit{ is an MD strategy}, w \in L^* >$

A polynomial algorithm to compute a basis of V_1

$V_1 \supseteq V_2 \supseteq V_3 \supseteq V_1$

Trace inequivalence problem for MDPs

$V_1 =< M_{\alpha_1}(w_1)M_{\alpha_2}(w_2)\ldots M_{\alpha_m}(w_m)\vec{\mathbf{1}}\colon \alpha_i \text{ is a memoryless strategy},\ w = w_1w_2\ldots w_m \in L^* >$

$V_2 =< M_\alpha(w)\vec{\mathbf{1}}\colon \alpha \text{ is a memoryless strategy}, w \in L^* >$

$\mu_1 \not\equiv \mu_2$ iff $\mu_1\vec{v} \neq \mu_2\vec{v}$ for some $\vec{v} \in V_2$

$V_3 =< M_\alpha(w)\vec{\mathbf{1}}\colon \alpha \text{ is an MD strategy}, w \in L^* >$

A polynomial algorithm to compute a basis of V_1

$$V_1 \supseteq V_2 \supseteq V_3 \supseteq V_1$$

$$V_1 = V_2 = V_3$$

Proof sketch

α_0 – any MD strategy
$\mathrm{A} = \{\alpha_0\} \cup \{\alpha_0^{s\to m}: s\ is\ a\ state, m\ is\ an\ action\ of\ s\}$
$B_0 = \{(\alpha_0, \varepsilon)\}$ – a set of vectors, e.g. $\vec{\mathbf{1}} = M_{\alpha_0}(\varepsilon)\vec{\mathbf{1}}$

Prove by induction
B_j is a basis for $V_1^j = < M_{\alpha_1}(w_1)M_{\alpha_2}(w_2) \dots M_{\alpha_k}(w_k)\vec{\mathbf{1}} \in \mathrm{V}_1: |w| \le j >$

In the (j+1)th iteration, if $\alpha \in A, a \in L, (\beta, w) \in B_j$ such that $M_\alpha(a)M_\beta(w)\vec{\mathbf{1}} \notin < B_{j+1} >$ add $(\alpha', aw)\ to\ B_{j+1}$, where

$$\alpha' = \begin{cases} \alpha(s): & if\ \overrightarrow{\boldsymbol{c_s}} \notin V_1^j \\ \beta(s): & otherwise \end{cases}$$

α' is an MD strategy

$V_1 = < M_{\alpha_1}(w_1) M_{\alpha_2}(w_2) \ldots M_{\alpha_m}(w_m)\vec{\mathbf{1}}: \alpha_i\ is\ a\ memoryless\ strategy,\ \ w = w_1 w_2 \ldots w_m \in L^* >$

$V_2 = < M_\alpha(w)\vec{\mathbf{1}}: \alpha\ is\ a\ memoryless\ strategy, w \in L^* >$

$\mu_1 \not\equiv \mu_2$ iff $\mu_1 \vec{v} \neq \mu_2 \vec{v}$ for some $\vec{v} \in V_2$

$V_3 = < M_\alpha(w)\vec{\mathbf{1}}: \alpha\ is\ an\ MD\ strategy, w \in L^* >$

A polynomial algorithm to compute

$B_n = \{(\alpha_0, \varepsilon), (\alpha_1, w_1), (\alpha_2, w_2), \ldots\}$ a basis of V_1 where each α_i is MD

$V_1 \supseteq V_2 \supseteq V_3 \supseteq V_1$

$V_1 = V_2 = V_3$

Trace inequivalence is easy!

Trace inequivalence problem for MDPs is in P (both memoryless and general strategies) [Kiefer and T., 2019]

Trace inequivalence is easy!

Trace inequivalence problem for MDPs is in P (both memoryless and general strategies) [Kiefer and T., 2019]

Trace equivalence problem (memoryless strategies) for MDPs is $\exists\mathbb{R}$-*complete* [Kiefer and T., 2019]

Trace equivalence problem (general strategies) for MDPs is undecidable [Fijalkow, Kiefer and Shirmohammadi, 2020]

Probabilistic bisimilarity inequivalence is easy!

Probabilistic bisimilarity inequivalence problem for MDPs is in P (both memoryless and general strategies) [Kiefer and T., 2019, 2022]

Probabilistic bisimilarity problem (memoryless strategies) for MDPs is NP-complete [Kiefer and T., 2019]

Probabilistic bisimilarity problem (general strategies) for MDPs is EXPTIME-complete [Kiefer and T., 2022]

Behavioural (in)equivalence problems

More notions:

$\varepsilon\text{-}bisimilarity, robust\ bisimilarity$, …

More models:

MDP-ST, POMDP, …

Behavioural (in)equivalence problems

More notions:

$\varepsilon\text{-}bisimilarity, robust\ bisimilarity$, …

More models:

MDP-ST, POMDP, …

Thank you!