

Computing Behavioral Distances, Compositionally

Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, Radu Mardare

Dept. of Computer Science, Aalborg University

MT-LAB Meeting
8-9 August, Copenhagen

Markov Decision Processes with Rewards

- + external nondeterminism + probabilistic behavior
- + many useful applications (A.I., planning, games, biology, ...)

Bisimilarity Distances

(bisimilarity is not robust: it only relates states with **identical** behaviors)

- + measure the behavioral similarity between states
- + support approximate reasoning on probabilistic systems
- + need of efficient methods for computing bisim. distances

Compositionality $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2 \otimes \dots \otimes \mathcal{M}_n$

- + may suffer from an exponential growth of the state space
(the parallel composition of n systems with m states has m^n states!)
- + exploit the structure of systems to compute bisim. distances

Markov Decision Processes with Rewards

- + external nondeterminism + probabilistic behavior
- + many useful applications (A.I., planning, games, biology, ...)

Bisimilarity Distances

(bisimilarity is not robust: it only relates states with **identical** behaviors)

- + measure the behavioral similarity between states
- + support approximate reasoning on probabilistic systems
- + need of efficient methods for computing bisim. distances

Compositionality $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2 \otimes \dots \otimes \mathcal{M}_n$

- + may suffer from an exponential growth of the state space
(the parallel composition of n systems with m states has m^n states!)
- + exploit the structure of systems to compute bisim. distances

Markov Decision Processes with Rewards

- + external nondeterminism + probabilistic behavior
- + many useful applications (A.I., planning, games, biology, ...)

Bisimilarity Distances

(bisimilarity is not robust: it only relates states with **identical** behaviors)

- + measure the behavioral similarity between states
- + support approximate reasoning on probabilistic systems
- + need of efficient methods for computing bisim. distances

Compositionality $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2 \otimes \dots \otimes \mathcal{M}_n$

- + may suffer from an exponential growth of the state space
(the parallel composition of n systems with m states has m^n states!)
- + exploit the structure of systems to compute bisim. distances

$$\mathcal{M} = (S, A, \tau, \rho)$$

Markov Decision Processes with Rewards (MDPs)

finite set of states

$$\mathcal{M} = (\mathcal{S}, A, \tau, \rho)$$

Markov Decision Processes with Rewards (MDPs)

The diagram shows the mathematical representation of a Markov Decision Process (MDP) as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \tau, \rho)$. Two yellow boxes with arrows point to the components: "finite set of states" points to \mathcal{S} , and "set of labels" points to \mathcal{A} .

$$\mathcal{M} = (\mathcal{S}, \mathcal{A}, \tau, \rho)$$

Markov Decision Processes with Rewards (MDPs)


finite set of states

set of labels



$$\mathcal{M} = (S, A, \tau, \rho)$$

probability transition function
 $\tau: S \times A \rightarrow \Delta(S)$

Markov Decision Processes with Rewards (MDPs)





Markov Decision Processes with Rewards (MDPs)





s_1 s_2 s_3



Markov Decision Processes with Rewards (MDPs)



Markov Decision Processes with Rewards (MDPs)





Markov Decision Processes with Rewards (MDPs)



Executions: $\omega = (s_0, a_0)(s_1, a_1) \dots$

Markov Decision Processes with Rewards (MDPs)





Executions: $\omega = (s_0, a_0)(s_1, a_1) \dots$

Discounted accumulated reward $\lambda \in (0, 1)$

$$R_\lambda(\omega) = \sum_{i \in \mathbb{N}} \lambda^i \cdot \rho(s_i, a_i)$$

Markov Decision Processes with Rewards (MDPs)



Executions: $\omega = (s_0, a_0)(s_1, a_1) \dots$

Discounted accumulated reward $\lambda \in (0, 1)$

$$R_\lambda(\omega) = \sum_{i \in \mathbb{N}} \lambda^i \cdot \rho(s_i, a_i)$$

Goal: To find policies $\pi: S \rightarrow A$ that maximize the expected value of R_λ on probabilistic executions starting from a given state.

Bisimilarity for MDPs

Extends probabilistic bisimilarity on Markov chains [Larsen-Skou'91]

Stochastic Bisimulation on \mathcal{M}

[Givan et al. AI'03]


Equivalence relation $R \subseteq S \times S$ such that,

$$s R t \implies \forall a \in A. \begin{cases} \rho(s, a) = \rho(t, a) \\ \forall R\text{-equiv. class } C. \sum_{u \in C} \tau(s, a)(c) = \sum_{u \in C} \tau(t, a)(c) \end{cases}$$


Stochastic Bisimilarity on \mathcal{M} :

$s \sim_{\mathcal{M}} t \iff s R t$ for some stochastic bisimulation R on \mathcal{M} .


Bisimilarity is not robust enough




Bisimilarity is not robust enough



Bisimilarity is not robust enough



Bisimilarity is not robust enough



From equivalences to distances

Pseudometrics $d: S \times S \rightarrow \mathbb{R}_{\geq 0}$ are the quantitative analogue of an equivalence relation

equivalence		pseudometric
$s \equiv s$	\rightsquigarrow	$d(s, s) = 0$
$s \equiv t \implies t \equiv s$	\rightsquigarrow	$d(s, t) = d(t, s)$
$s \cong u \wedge u \cong t \implies s \cong t$	\rightsquigarrow	$d(s, u) + d(u, t) \geq d(s, t)$

Bisimilarity Pseudometric: $d(s, t) = 0 \iff s \sim t$

From equivalences to distances

Pseudometrics $d: S \times S \rightarrow \mathbb{R}_{\geq 0}$ are the quantitative analogue of an equivalence relation


equivalence		pseudometric
$s \equiv s$	\rightsquigarrow	$d(s, s) = 0$
$s \equiv t \implies t \equiv s$	\rightsquigarrow	$d(s, t) = d(t, s)$
$s \cong u \wedge u \cong t \implies s \cong t$	\rightsquigarrow	$d(s, u) + d(u, t) \geq d(s, t)$

Bisimilarity Pseudometric: $d(s, t) = 0 \iff s \sim t$


We consider the λ -discounted bisimilarity distances $\delta_\lambda: S \times S \rightarrow \mathbb{R}_{\geq 0}$ proposed by Ferns et al. [UAI'04]

From equivalences to distances

Bisimulation



Metric analogue




Kantorovich Metric: $\mathcal{T}_d: \Delta(S) \times \Delta(S) \rightarrow \mathbb{R}_{\geq 0}$

The distance between $\tau(s, a)$ and $\tau(t, a)$
is the optimal value of a **Transportation Problem**

$$\mathcal{T}_d(\tau(s, a), \tau(t, a)) = \min \left\{ \sum_{u, v \in S} d(u, v) \cdot \omega(u, v) \mid \begin{array}{l} \forall u \in S \sum_{v \in S} \omega(u, v) = \tau(s, a)(u) \\ \forall v \in S \sum_{u \in S} \omega(u, v) = \tau(t, a)(v) \end{array} \right\}$$

ω can be understood as **transportation** of $\pi(s, a)$ to $\pi(t, a)$.




Kantorovich Metric: $\mathcal{T}_d: \Delta(S) \times \Delta(S) \rightarrow \mathbb{R}_{\geq 0}$

The distance between $\tau(s, a)$ and $\tau(t, a)$
is the optimal value of a **Transportation Problem**

$$\mathcal{T}_d(\tau(s, a), \tau(t, a)) = \min \left\{ \sum_{u, v \in S} d(u, v) \cdot \omega(u, v) \mid \begin{array}{l} \forall u \in S \sum_{v \in S} \omega(u, v) = \tau(s, a)(u) \\ \forall v \in S \sum_{u \in S} \omega(u, v) = \tau(t, a)(v) \end{array} \right\}$$

matching $\omega \in \Pi(\tau(s, a), \tau(t, a))$

ω can be understood as **transportation** of $\pi(s, a)$ to $\pi(t, a)$.



The bisimilarity pseudometric $\delta_\lambda^{\mathcal{M}}$ is the **least fixed point** of the following operator on pseudometrics

$$F_\lambda^{\mathcal{M}}(d)(s, t) = \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \mathcal{T}_d(\tau(s, a), \tau(t, a)) \right\}$$

The bisimilarity pseudometric $\delta_\lambda^{\mathcal{M}}$ is the **least fixed point** of the following operator on pseudometrics

$$F_\lambda^{\mathcal{M}}(d)(s, t) = \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \mathcal{T}_d(\tau(s, a), \tau(t, a)) \right\}$$

distance between rewards

The bisimilarity pseudometric $\delta_\lambda^{\mathcal{M}}$ is the **least fixed point** of the following operator on pseudometrics

$$F_\lambda^{\mathcal{M}}(d)(s, t) = \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \mathcal{T}_d(\tau(s, a), \tau(t, a)) \right\}$$

distance between rewards

and recursively...

distance between
transition probabilities

Algebraic operators on MDPs

Systems can be conveniently represented as the algebraic composition of simpler sub-systems

Algebraic operators on MDPs

Systems can be conveniently represented as the algebraic composition of simpler sub-systems

How to define operators on MDPs?

$$\mathcal{M}_1 \otimes \mathcal{M}_2 = (\underset{\substack{\uparrow \\ \text{set of} \\ \text{states}}}{S} , \underset{\substack{\uparrow \\ \text{set of} \\ \text{actions}}}{A} , \underset{\substack{\uparrow \\ \text{probability} \\ \text{transition} \\ \text{function}}}{\tau} , \underset{\substack{\uparrow \\ \text{reward} \\ \text{function}}}{\rho})$$

Algebraic operators on MDPs

Systems can be conveniently represented as the algebraic composition of simpler sub-systems

How to define operators on MDPs?

$$\mathcal{M}_1 \otimes \mathcal{M}_2 = (\mathcal{S}_1 \times \mathcal{S}_2, \mathcal{A}_1 \otimes_A \mathcal{A}_2, \tau_1 \otimes_\tau \tau_2, \rho_1 \otimes_\rho \tau_2)$$

set of
states

set of
actions

probability
transition
function

reward
function

Example 1: Synchronous parallel composition

$$\mathcal{M}_1 \mid \mathcal{M}_2 = (\mathcal{S}_1 \times \mathcal{S}_2, A_1 \cap A_2, \tau_1 \mid_{\tau} \tau_2, \rho_1 \mid_{\rho} \rho_2)$$

$$\frac{s_1 \xrightarrow{a[r_1]}_{p_1} s'_1 \quad s_2 \xrightarrow{a[r_2]}_{p_2} s'_2}{s_1 \mid s_2 \xrightarrow{a[r_1+r_2]}_{p_1 \cdot p_2} s'_1 \mid s'_2}$$

Example 1: Synchronous parallel composition

$$\mathcal{M}_1 \mid \mathcal{M}_2 = (\mathcal{S}_1 \times \mathcal{S}_2, \mathcal{A}_1 \cap \mathcal{A}_2, \tau_1 \mid_{\tau} \tau_2, \rho_1 \mid_{\rho} \rho_2)$$

$$\frac{s_1 \xrightarrow{a[r_1]}_{\rho_1} s'_1 \quad s_2 \xrightarrow{a[r_2]}_{\rho_2} s'_2}{s_1 \mid s_2 \xrightarrow{a[r_1+r_2]}_{\rho_1 \cdot \rho_2} s'_1 \mid s'_2}$$

$$\begin{aligned}(\tau_1 \mid_{\tau} \tau_2)((s_1, s_2), a)(s'_1, s'_2) &= \tau_1(s_1, a)(s'_1) \cdot \tau_2(s_2, a)(s'_2), \\(\rho_1 \mid_{\rho} \rho_2)((s_1, s_2), a) &= \rho_1(s_1, a) + \rho_2(s_2, a).\end{aligned}$$

Example 2: CCS-like parallel composition

$$\mathcal{M}_1 \parallel \mathcal{M}_2 = (\mathcal{S}_1 \times \mathcal{S}_2, A_1 \cup A_2, \tau_1 \parallel_{\tau} \tau_2, \rho_1 \parallel_{\rho} \rho_2)$$

$$\frac{s_1 \xrightarrow{a[r]}_p s'_1 \quad a \notin A_2}{s_1 \parallel s_2 \xrightarrow{a[r]}_p s'_1 \parallel s_2}$$

$$\frac{s_2 \xrightarrow{a[r]}_p s'_2 \quad a \notin A_1}{s_1 \parallel s_2 \xrightarrow{a[r]}_p s_1 \parallel s'_2}$$

$$\frac{s_1 \xrightarrow{a[r_1]}_{p_1} s'_1 \quad s_2 \xrightarrow{a[r_2]}_{p_2} s'_2}{s_1 \parallel s_2 \xrightarrow{a[r_1+r_2]}_{p_1 \cdot p_2} s'_1 \parallel s'_2}$$

Example 2: CCS-like parallel composition

$$\mathcal{M}_1 \parallel \mathcal{M}_2 = (S_1 \times S_2, A_1 \cup A_2, \tau_1 \parallel_{\tau} \tau_2, \rho_1 \parallel_{\rho} \rho_2)$$

$$\frac{s_1 \xrightarrow{a[r]}_p s'_1 \quad a \notin A_2}{s_1 \parallel s_2 \xrightarrow{a[r]}_p s'_1 \parallel s_2} \qquad \frac{s_2 \xrightarrow{a[r]}_p s'_2 \quad a \notin A_1}{s_1 \parallel s_2 \xrightarrow{a[r]}_p s_1 \parallel s'_2}$$

$$\frac{s_1 \xrightarrow{a[r_1]}_{p_1} s'_1 \quad s_2 \xrightarrow{a[r_2]}_{p_2} s'_2}{s_1 \parallel s_2 \xrightarrow{a[r_1+r_2]}_{p_1 \cdot p_2} s'_1 \parallel s'_2}$$

$$(\tau_1 \parallel_{\tau} \tau_2)((s_1, s_2), a)(s'_1, s'_2) = \begin{cases} \tau_1(s_1, a)(s'_1) & \text{if } a \notin A_2 \text{ and } s_2 = s'_2 \\ \tau_2(s_2, a)(s'_2) & \text{if } a \notin A_1 \text{ and } s_1 = s'_1 \\ \tau_1(s_1, a)(s'_1) \cdot \tau_2(s_2, a)(s'_2) & \text{if } a \in A_1 \cap A_2 \\ 0 & \text{otherwise} \end{cases}$$

$$(\rho_1 \parallel_{\rho} \rho_2)((s_1, s_2), a) = \begin{cases} \rho_1(s_1, a) & \text{if } a \notin A_2 \\ \rho_2(s_2, a) & \text{if } a \notin A_1 \\ \rho_1(s_1, a) + \rho_2(s_2, a) & \text{if } a \in A_1 \cap A_2 \end{cases}$$

Metric analogue of congruence

Operators over MDPs are well-behaved when they are congruential w.r.t. bisimilarity:

$$s_1 \sim_{\mathcal{M}_1} t_1 \text{ and } s_2 \sim_{\mathcal{M}_2} t_2 \implies s_1 \otimes s_2 \sim_{\mathcal{M}_1 \otimes \mathcal{M}_2} t_1 \otimes t_2$$

Metric analogue of congruence

Operators over MDPs are well-behaved when they are congruential w.r.t. bisimilarity:

$$s_1 \sim_{\mathcal{M}_1} t_1 \text{ and } s_2 \sim_{\mathcal{M}_2} t_2 \implies s_1 \otimes s_2 \sim_{\mathcal{M}_1 \otimes \mathcal{M}_2} t_1 \otimes t_2$$

What is the quantitative analogue of congruence?

Metric analogue of congruence

Operators over MDPs are well-behaved when they are congruential w.r.t. bisimilarity:

$$s_1 \sim_{\mathcal{M}_1} t_1 \text{ and } s_2 \sim_{\mathcal{M}_2} t_2 \implies s_1 \otimes s_2 \sim_{\mathcal{M}_1 \otimes \mathcal{M}_2} t_1 \otimes t_2$$

What is the quantitative analogue of congruence?

$$+ \left. \begin{array}{l} \delta_\lambda^{\mathcal{M}_1}(s_1, t_1) = 0 \\ \delta_\lambda^{\mathcal{M}_2}(s_2, t_2) = 0 \end{array} \right\} \implies \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2}(s_1 \otimes s_2, t_1 \otimes t_2) = 0$$

Metric analogue of congruence

Operators over MDPs are well-behaved when they are congruential w.r.t. bisimilarity:

$$s_1 \sim_{\mathcal{M}_1} t_1 \text{ and } s_2 \sim_{\mathcal{M}_2} t_2 \implies s_1 \otimes s_2 \sim_{\mathcal{M}_1 \otimes \mathcal{M}_2} t_1 \otimes t_2$$

What is the quantitative analogue of congruence?

$$+ \left. \begin{array}{l} \delta_\lambda^{\mathcal{M}_1}(s_1, t_1) = 0 \\ \delta_\lambda^{\mathcal{M}_2}(s_2, t_2) = 0 \end{array} \right\} \implies \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2}(s_1 \otimes s_2, t_1 \otimes t_2) = 0$$

$$+ \delta_\lambda^{\mathcal{M}_1}(s_1, t_1) + \delta_\lambda^{\mathcal{M}_2}(s_2, t_2) \geq \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2}(s_1 \otimes s_2, t_1 \otimes t_2)$$

Metric analogue of congruence

Operators over MDPs are well-behaved when they are congruential w.r.t. bisimilarity:

$$s_1 \sim_{\mathcal{M}_1} t_1 \text{ and } s_2 \sim_{\mathcal{M}_2} t_2 \implies s_1 \otimes s_2 \sim_{\mathcal{M}_1 \otimes \mathcal{M}_2} t_1 \otimes t_2$$

What is the quantitative analogue of congruence?

$$+ \left. \begin{array}{l} \delta_\lambda^{\mathcal{M}_1}(s_1, t_1) = 0 \\ \delta_\lambda^{\mathcal{M}_2}(s_2, t_2) = 0 \end{array} \right\} \implies \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2}(s_1 \otimes s_2, t_1 \otimes t_2) = 0$$

$$+ \delta_\lambda^{\mathcal{M}_1}(s_1, t_1) + \delta_\lambda^{\mathcal{M}_2}(s_2, t_2) \geq \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2}(s_1 \otimes s_2, t_1 \otimes t_2)$$

$$+ \|\delta_\lambda^{\mathcal{M}_1}, \delta_\lambda^{\mathcal{M}_2}\|_1 \sqsupseteq \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2} \quad (\otimes \text{ is non-extensive})$$

Metric analogue of congruence

Operators over MDPs are well-behaved when they are congruential w.r.t. bisimilarity:

$$s_1 \sim_{\mathcal{M}_1} t_1 \text{ and } s_2 \sim_{\mathcal{M}_2} t_2 \implies s_1 \otimes s_2 \sim_{\mathcal{M}_1 \otimes \mathcal{M}_2} t_1 \otimes t_2$$

What is the quantitative analogue of congruence?

$$+ \left. \begin{array}{l} \delta_\lambda^{\mathcal{M}_1}(s_1, t_1) = 0 \\ \delta_\lambda^{\mathcal{M}_2}(s_2, t_2) = 0 \end{array} \right\} \implies \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2}(s_1 \otimes s_2, t_1 \otimes t_2) = 0$$

$$+ \delta_\lambda^{\mathcal{M}_1}(s_1, t_1) + \delta_\lambda^{\mathcal{M}_2}(s_2, t_2) \geq \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2}(s_1 \otimes s_2, t_1 \otimes t_2)$$

$$+ \|\delta_\lambda^{\mathcal{M}_1}, \delta_\lambda^{\mathcal{M}_2}\|_p \sqsupseteq \delta_\lambda^{\mathcal{M}_1 \otimes \mathcal{M}_2} \quad (\otimes \text{ is } p\text{-non-extensive})$$

Safe algebraic operators on MDPs

We characterized a class of operators on MDPs

p -Safe operators

$$F_{\lambda}^{\mathcal{M}_1 \otimes \mathcal{M}_2}(\|d_1, d_2\|_p) \subseteq \|F_{\lambda}^{\mathcal{M}_1}(d_1), F_{\lambda}^{\mathcal{M}_2}(d_2)\|_p$$

Theorem: p -Safeness \implies non-extensiveness

Checking p -Safeness is simpler than checking non-extensiveness

($\delta_{\lambda}^{\mathcal{M}}$ is defined as the least fixed point of $F_{\lambda}^{\mathcal{M}}$)

Safe algebraic operators on MDPs

We characterized a class of operators on MDPs

p -Safe operators

$$F_{\lambda}^{\mathcal{M}_1 \otimes \mathcal{M}_2}(\|d_1, d_2\|_p) \subseteq \|F_{\lambda}^{\mathcal{M}_1}(d_1), F_{\lambda}^{\mathcal{M}_2}(d_2)\|_p$$

Theorem: p -Safeness \implies non-extensiveness

Checking p -Safeness is simpler than checking non-extensiveness
($\delta_{\lambda}^{\mathcal{M}}$ is defined as the least fixed point of $F_{\lambda}^{\mathcal{M}}$)

✓ Synch. parallel comp.

✓ CCS-like parallel comp.

Computing the behavioral distance

given $s, t \in S$, to compute $\delta_\lambda^{\mathcal{M}}(s, t)$

On-the-fly algorithm

[Bacci², Larsen, Mardare TACAS'13]

- + lazy exploration of \mathcal{M}
- + save comput. time + space

Compositional strategy

- + exploit the compositional structure of $\mathcal{M}_1 \otimes \mathcal{M}_2$

Alternative characterization of $\delta_\lambda^{\mathcal{M}}$

$$F_\lambda^{\mathcal{M}}(d)(s, t) = \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \mathcal{T}_d(\tau(s, a), \tau(t, a)) \right\}$$

Alternative characterization of δ_λ^M

$$\begin{aligned} F_\lambda^M(d)(s, t) &= \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \mathcal{T}_d(\tau(s, a), \tau(t, a)) \right\} \\ &= \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \min_{\omega \in \Pi(\tau(s, a), \tau(t, a))} \sum_{u, v \in S} d(u, v) \cdot \omega(u, v) \right\} \end{aligned}$$

Alternative characterization of $\delta_\lambda^{\mathcal{M}}$

$$\begin{aligned} F_\lambda^{\mathcal{M}}(d)(s, t) &= \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \mathcal{T}_d(\tau(s, a), \tau(t, a)) \right\} \\ &= \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \cdot \min_{\omega \in \Pi(\tau(s, a), \tau(t, a))} \sum_{u, v \in S} d(u, v) \cdot \omega(u, v) \right\} \end{aligned}$$

Coupling: $\mathcal{C} = \left\{ \omega_{s,t}^a \in \Pi(\tau(s, a), \tau(t, a)) \right\}_{s,t \in S}^{a \in A}$


$$\Gamma_\lambda^{\mathcal{C}}(d)(s, t) = \max_{a \in A} \left\{ |\rho(s, a) - \rho(t, a)| + \lambda \sum_{u, v \in S} d(u, v) \cdot \omega_{s,t}^a(u, v) \right\}$$

we call **discrepancy**, $\gamma_\lambda^{\mathcal{C}}$, the least fixed point of $\Gamma_\lambda^{\mathcal{C}}$


Theorem:

$$\delta_\lambda^{\mathcal{M}} = \min \{ \gamma_\lambda^{\mathcal{C}} \mid \mathcal{C} \text{ coupling for } \mathcal{M} \} \text{ for all } \lambda \in (0, 1).$$


$$C_1 \triangleleft_{\lambda} C_2 \iff \gamma_{\lambda}^{C_1} \sqsubseteq \gamma_{\lambda}^{C_2}$$



$$C_1 \triangleleft_{\lambda} C_2 \iff \gamma_{\lambda}^{C_1} \sqsubseteq \gamma_{\lambda}^{C_2}$$



$$C_1 \triangleleft_{\lambda} C_2 \iff \gamma_{\lambda}^{C_1} \sqsubseteq \gamma_{\lambda}^{C_2}$$



Greedy strategy

Moving Criterion:

$$C_i = \{\dots, \omega_{u,v}^a, \dots\}$$


$\omega_{u,v}^a$ not opt. w.r.t. $TP(\gamma_{\lambda}^{C_i}, \tau(u, a), \tau(v, a))$

Improvement:

$C_{i+1} = \{\dots, \omega^*, \dots\}$, where

ω^* optimal sol. for $TP(\gamma_{\lambda}^{C_i}, \tau(u, a), \tau(v, a))$

$$C_1 \triangleleft_{\lambda} C_2 \iff \gamma_{\lambda}^{C_1} \sqsubseteq \gamma_{\lambda}^{C_2}$$



Greedy strategy

Moving Criterion:

$$C_i = \{\dots, \omega_{u,v}^a, \dots\}$$

$\omega_{u,v}^a$ not opt. w.r.t. $TP(\gamma_{\lambda}^{C_i}, \tau(u, a), \tau(v, a))$

Improvement:


$C_{i+1} = \{\dots, \omega^*, \dots\}$, where

ω^* optimal sol. for $TP(\gamma_{\lambda}^{C_i}, \tau(u, a), \tau(v, a))$

Theorem

- + each step ensures $C_{i+1} \triangleleft_{\lambda} C_i$
- + moving criterion holds until $\gamma_{\lambda}^{C_i} \neq \delta_{\lambda}$
- + the method always terminates

$$C_1 \triangleleft_{\lambda} C_2 \iff \gamma_{\lambda}^{C_1} \sqsubseteq \gamma_{\lambda}^{C_2}$$



Greedy strategy

Moving Criterion:

$$C_i = \{\dots, \omega_{u,v}^a, \dots\}$$

$\omega_{u,v}^a$ not opt. w.r.t. $TP(\gamma_{\lambda}^{C_i}, \tau(u, a), \tau(v, a))$

Improvement:

$C_{i+1} = \{\dots, \omega^*, \dots\}$, where

ω^* optimal sol. for $TP(\gamma_{\lambda}^{C_i}, \tau(u, a), \tau(v, a))$

Theorem

- + each step ensures $C_{i+1} \triangleleft_{\lambda} C_i$
- + moving criterion holds until $\gamma_{\lambda}^{C_i} \neq \delta_{\lambda}$
- + the method always terminates

A Compositional Heuristic

Let $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2$ and \otimes be non-extensive, then

$$\delta_\lambda^{\mathcal{M}} \sqsubseteq \|\delta_\lambda^{\mathcal{M}_1}, \delta_\lambda^{\mathcal{M}_2}\|_p$$

A Compositional Heuristic

Let $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2$ and \otimes be non-extensive, then

$$\begin{array}{ccc} \delta_\lambda^{\mathcal{M}} & \subseteq & \|\delta_\lambda^{\mathcal{M}_1}, \delta_\lambda^{\mathcal{M}_2}\|_p \\ // & & // \\ \gamma_\lambda^{\mathcal{D}} & & \|\gamma_\lambda^{\mathcal{D}_1}, \gamma_\lambda^{\mathcal{D}_2}\|_p \end{array}$$

A Compositional Heuristic

Let $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2$ and \otimes be non-extensive, then

$$\begin{aligned} \delta_\lambda^{\mathcal{M}} &\subseteq \|\delta_\lambda^{\mathcal{M}_1}, \delta_\lambda^{\mathcal{M}_2}\|_p \\ // & \qquad \qquad // \\ \gamma_\lambda^{\mathcal{D}} &\subseteq \gamma_\lambda^{\mathcal{D}^*} \subseteq \|\gamma_\lambda^{\mathcal{D}_1}, \gamma_\lambda^{\mathcal{D}_2}\|_p \end{aligned}$$

A good starting coupling should not exceed the upper-bound given by non-extensiveness!

A Compositional Heuristic

Let $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2$ and \otimes be non-extensive, then

$$\begin{array}{ccc} \delta_\lambda^{\mathcal{M}} & \subseteq & \|\delta_\lambda^{\mathcal{M}_1}, \delta_\lambda^{\mathcal{M}_2}\|_p \\ // & & // \\ \gamma_\lambda^{\mathcal{D}} & \subseteq & \gamma_\lambda^{\mathcal{D}^*} \subseteq \|\gamma_\lambda^{\mathcal{D}_1}, \gamma_\lambda^{\mathcal{D}_2}\|_p \end{array}$$

A good starting coupling should not exceed the upper-bound given by non-extensiveness!

Remark: \mathcal{D}^* should be obtained from \mathcal{D}_1 and \mathcal{D}_2

Lifting algebraic operators on Couplings

Lifting operator

$$\begin{array}{ccc} \mathcal{M}_1, & \mathcal{M}_2 \vdash & \mathcal{M}_1 \otimes \mathcal{M}_2 \\ \downarrow & \downarrow & \downarrow \\ \mathcal{C}_1, & \mathcal{C}_2 \vdash & \mathcal{C}_1 \otimes^* \mathcal{C}_2 \end{array}$$

Lifting algebraic operators on Couplings

Lifting operator

$$\begin{array}{ccc} \mathcal{M}_1, & \mathcal{M}_2 \mapsto & \mathcal{M}_1 \otimes \mathcal{M}_2 \\ \downarrow & \downarrow & \downarrow \\ \mathcal{C}_1, & \mathcal{C}_2 \mapsto & \mathcal{C}_1 \otimes^* \mathcal{C}_2 \end{array}$$

+

p-Safe lifting operator

$$\Gamma_{\lambda}^{\mathcal{C}_1 \otimes^* \mathcal{C}_2}(\|d_1, d_2\|_p) \subseteq \|\Gamma_{\lambda}^{\mathcal{C}_1}(d_1), \Gamma_{\lambda}^{\mathcal{C}_2}(d_2)\|_p$$

=

Lifting algebraic operators on Couplings

Lifting operator

$$\begin{array}{ccc} \mathcal{M}_1, & \mathcal{M}_2 & \vdash \mathcal{M}_1 \otimes \mathcal{M}_2 \\ \vdots & \vdots & \vdots \\ \mathcal{C}_1, & \mathcal{C}_2 & \vdash \mathcal{C}_1 \otimes^* \mathcal{C}_2 \end{array}$$

+

p-Safe lifting operator



$$\Gamma_{\lambda}^{\mathcal{C}_1 \otimes^* \mathcal{C}_2}(\|d_1, d_2\|_p) \sqsubseteq \|\Gamma_{\lambda}^{\mathcal{C}_1}(d_1), \Gamma_{\lambda}^{\mathcal{C}_1}(d_2)\|_p$$

=



$$\delta_{\lambda}^{\mathcal{M}_1 \otimes \mathcal{M}_2} \sqsubseteq \gamma_{\lambda}^{\mathcal{D}_1 \otimes^* \mathcal{D}_2} \sqsubseteq \|\delta_{\lambda}^{\mathcal{M}_1}, \delta_{\lambda}^{\mathcal{M}_2}\|_p$$

where \mathcal{D}_i is a coupling for \mathcal{M}_i minimal w.r.t. \triangleleft_{λ}

The Pipeline Example



The Pipeline Example



Experimental Results

Query	Instance	OTF	COTF	# States
All pairs	$E_0 \parallel E_1$	0.654791	0.97248	9
	$E_1 \parallel E_2$	0.702105	0.801121	9
	$E_0 \parallel E_0 \parallel E_1$	48.5982	13.5731	27
	$E_0 \parallel E_1 \parallel E_2$	23.1984	19.9137	27
	$E_0 \parallel E_1 \parallel E_1$	126.335	13.6483	27
	$E_0 \parallel E_0 \parallel E_0$	49.1167	14.1075	27
Single pair	$E_0 \parallel E_0 \parallel E_0 \parallel E_1 \parallel E_1$	16.7027	11.6919	243
	$E_0 \parallel E_1 \parallel E_0 \parallel E_1 \parallel E_1$	20.2666	16.6274	243
	$E_2 \parallel E_1 \parallel E_0 \parallel E_1 \parallel E_1$	22.8357	10.4844	243
	$E_1 \parallel E_2 \parallel E_0 \parallel E_0 \parallel E_2$	11.7968	6.76188	243
	$E_1 \parallel E_2 \parallel E_0 \parallel E_0 \parallel E_2 \parallel E_2$	Time-out	79.902	729

Results

- + generic definition of algebraic operators on MDPs
- + characterized a well-behaved class of operators (p -Safeness)
- + on-the-fly algorithm for behavioral pseudometrics
 - + exact
 - + avoids entire exploration of the state space
 - + exploit compositional structure of the model (**first proposal!**)
- + developed a proof of concept prototype
- + performs, on average, better than other proposals

Future work

- + beyond non-extensiveness (continuous operators)
- + formal analysis of time/space complexity
- + apply similar techniques on CTMCs, CTMDPs, etc. . .