Bisimulation Metrics for Weighted Automata

B. Balle[1,*], *P. Gourdeau*[2] and P. Panangaden[2]

[1]Amazon, Cambridge (UK)

[*](work done while at Lancaster University)

[2]McGill University, Montreal (Canada)

ICALP 2017

Motivation

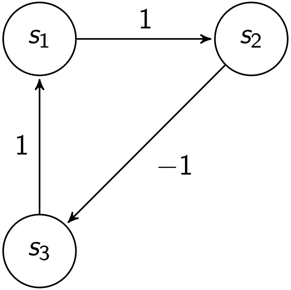

$=$

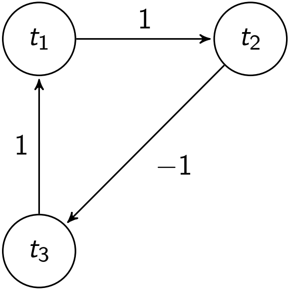

Motivation

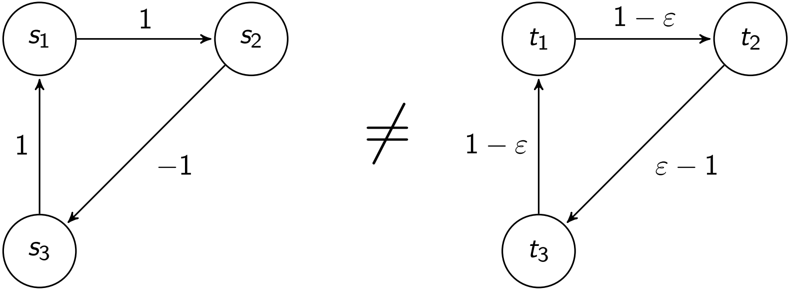

Motivation

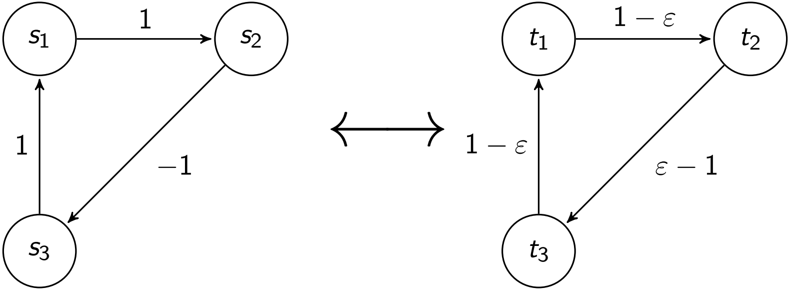

s_1
1
s_2
1
-1
s_3
$\longleftrightarrow$
t_1
$1-\varepsilon$
t_2
$1-\varepsilon$
$\varepsilon-1$
t_3

Contributions

1. The construction of bisimulation seminorms and the associated pseudometrics on WFA. We exhibit them in closed-form.

Contributions

1. The construction of bisimulation seminorms and the associated pseudometrics on WFA. We exhibit them in closed-form.
2. We obtain metrics on the space of weighted languages from the metrics on WFA.

Contributions

1. The construction of bisimulation seminorms and the associated pseudometrics on WFA. We exhibit them in closed-form.
2. We obtain metrics on the space of weighted languages from the metrics on WFA.
3. We show two continuity properties of the metric; one using definitions due to Jaeger et al. and the other developed here.

Contributions

1. The construction of bisimulation seminorms and the associated pseudometrics on WFA. We exhibit them in closed-form.
2. We obtain metrics on the space of weighted languages from the metrics on WFA.
3. We show two continuity properties of the metric; one using definitions due to Jaeger et al. and the other developed here.
4. We show undecidability results for computing our metrics.

Contributions

1. The construction of bisimulation seminorms and the associated pseudometrics on WFA. We exhibit them in closed-form.
2. We obtain metrics on the space of weighted languages from the metrics on WFA.
3. We show two continuity properties of the metric; one using definitions due to Jaeger et al. and the other developed here.
4. We show undecidability results for computing our metrics.
5. Nevertheless, we show that one can successfully exploit these metrics for applications in spectral learning.

Contributions

1. The construction of bisimulation seminorms and the associated pseudometrics on WFA. We exhibit them in closed-form.
2. We obtain metrics on the space of weighted languages from the metrics on WFA.
3. We show two continuity properties of the metric; one using definitions due to Jaeger et al. and the other developed here.
4. We show undecidability results for computing our metrics.
5. Nevertheless, we show that one can successfully exploit these metrics for applications in spectral learning.
6. Bisimulation for pseudometrics were first defined for LMPs in 1999 by Desharnais, Gupta, Jagadeesan and Panangaden and have been studied and developed for other models since then.

Weighted Finite Automata

Definition

A *weighted finite automaton* (WFA) is a tuple $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ where

Weighted Finite Automata

Definition

A *weighted finite automaton* (WFA) is a tuple $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ where

- Σ is a finite alphabet,
- V is a finite-dimensional vector space,
- $\alpha \in V$ is a vector representing the initial weights,
- $\beta \in V^*$ is a linear form representing the final weights,
- $\tau_\sigma : V \to V$ is a linear map representing the transition indexed by $\sigma \in \Sigma$.

Weighted Finite Automata

Definition

A *weighted finite automaton* (WFA) is a tuple $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma\in\Sigma} \rangle$ where

- Σ is a finite alphabet,
- V is a finite-dimensional vector space,
- $\alpha \in V$ is a vector representing the initial weights,
- $\beta \in V^*$ is a linear form representing the final weights,
- $\tau_\sigma : V \to V$ is a linear map representing the transition indexed by $\sigma \in \Sigma$.

Given a word $x = x_1 \dots x_n \in \Sigma^*$, the automaton $\mathcal{A}$ realizes the function $f_{\mathcal{A}} : \Sigma^\star \to \mathbb{R}$ defined by

$$f_{\mathcal{A}}(x) = \beta(\tau_{x_n}(\dots \tau_{x_1}(\alpha))) = \beta(\tau_x(\alpha)) \ .$$

Weighted Finite Automata

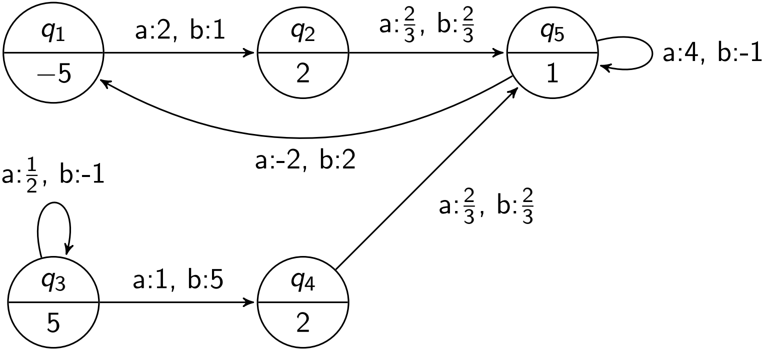

Figure 1: An example of a WFA where $V = \mathbb{R}^5$ (with the standard basis) and $\Sigma = \{a, b\}$. The final weights of the states are in the lower half of the circles and the initial weights are omitted.

Bisimulation (Boreale)

Definition

A *linear bisimulation* for a weighted automaton $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ on a vector space V is a linear subspace $W \subseteq V$ satisfying:

Bisimulation (Boreale)

Definition

A *linear bisimulation* for a weighted automaton $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ on a vector space V is a linear subspace $W \subseteq V$ satisfying:

1. $W \subseteq \ker(\beta)$, and
2. $\tau_\sigma(W) \subseteq W$ for all $\sigma \in \Sigma$.

Bisimulation (Boreale)

Definition

A *linear bisimulation* for a weighted automaton $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ on a vector space V is a linear subspace $W \subseteq V$ satisfying:

1. $W \subseteq \ker(\beta)$, and
2. $\tau_\sigma(W) \subseteq W$ for all $\sigma \in \Sigma$.

Two states $u, v \in V$ are called *W-bisimilar* if $u - v \in W$. This is denoted $u \sim_W v$

Bisimulation

Definition

For a given automaton $\mathcal{A}$, two states $u, v \in V$ are called $\mathcal{A}$-*bisimilar* if there exists a linear bisimulation $W \subseteq V$ such that $u - v \in W$.

Bisimulation

Definition

For a given automaton $\mathcal{A}$, two states $u, v \in V$ are called $\mathcal{A}$-*bisimilar* if there exists a linear bisimulation $W \subseteq V$ such that $u - v \in W$.

Theorem (Boreale)

For every WFA $\mathcal{A}$ there exists a largest linear bisimulation $W_{\mathcal{A}}$ for $\mathcal{A}$ such that $\sim_{\mathcal{A}} \equiv \sim_{W_{\mathcal{A}}}$.

Bisimulation

Definition

For a given automaton $\mathcal{A}$, two states $u, v \in V$ are called $\mathcal{A}$-*bisimilar* if there exists a linear bisimulation $W \subseteq V$ such that $u - v \in W$.

Theorem (Boreale)

For every WFA $\mathcal{A}$ there exists a largest linear bisimulation $W_{\mathcal{A}}$ for $\mathcal{A}$ such that $\sim_{\mathcal{A}} \equiv \sim_{W_{\mathcal{A}}}$.

We want a *quantitative analogue* of this relation.

Bisimulation

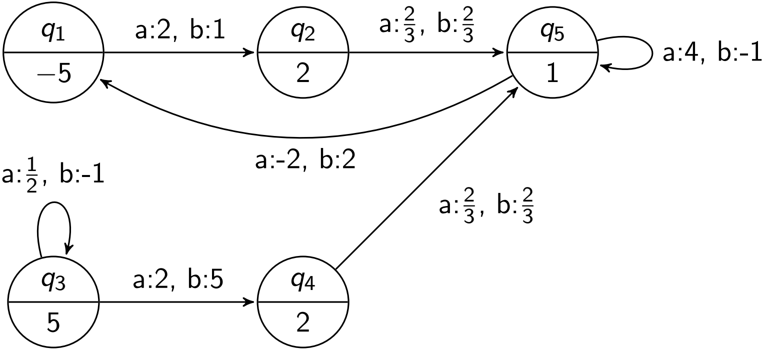

Figure 2: Here a linear bisimulation would be $W = \{(0\ \lambda\ 0\ \ -\lambda\ 0)^T : \lambda \in \mathbb{R}\}$.

Framework

S: set of all seminorms on V.

Framework

S: set of all seminorms on V.

WFA $\mathcal{A}$ with

$\gamma > 0$ and $\rho(\mathcal{A})\gamma < 1$

$F_{\mathcal{A},\gamma} : S \to S$

$F_{\mathcal{A},\gamma}(s)(v) = |\beta(v)| + \gamma \max_{\sigma\in\Sigma} s(\tau_\sigma(v))$

$\rho(\mathcal{A})$: joint spectral radius of the transition maps $\{\tau_\sigma\}_{\sigma\in\Sigma}$.

Framework

S: set of all seminorms on V.

WFA $\mathcal{A}$ with

$\gamma > 0$ and $\rho(\mathcal{A})\gamma < 1$

$F_{\mathcal{A},\gamma} : S \to S$

$F_{\mathcal{A},\gamma}(s)(v) = |\beta(v)| + \gamma \max_{\sigma \in \Sigma} s(\tau_\sigma(v))$

$F_{\mathcal{A},\gamma}$ is contractive.

$\rho(\mathcal{A})$: joint spectral radius of the transition maps $\{\tau_\sigma\}_{\sigma \in \Sigma}$.

S: set of all seminorms on V.

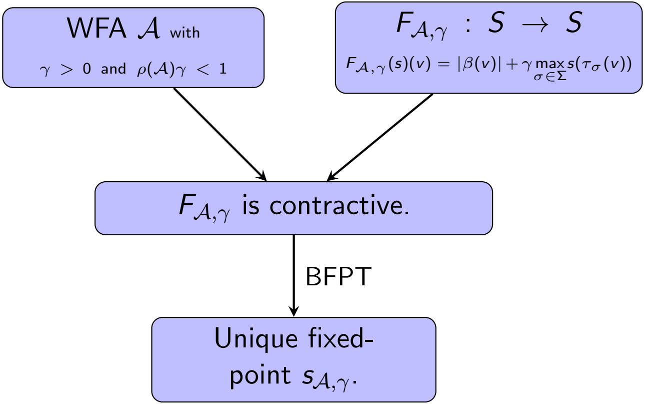

$\rho(\mathcal{A})$: joint spectral radius of the transition maps $\{\tau_\sigma\}_{\sigma \in \Sigma}$.

Framework

S: set of all seminorms on V.

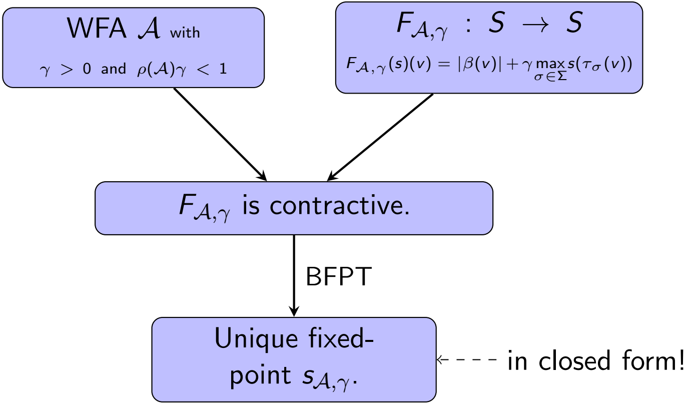

$\rho(\mathcal{A})$: joint spectral radius of the transition maps $\{\tau_\sigma\}_{\sigma \in \Sigma}$.

Framework

S: set of all seminorms on V.

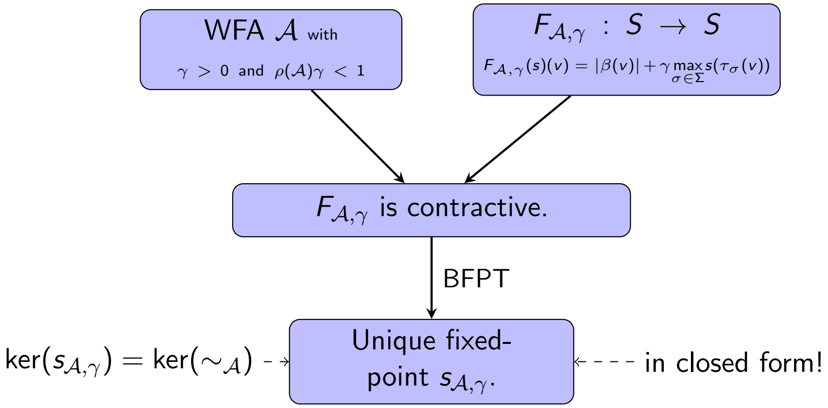

$\rho(\mathcal{A})$: joint spectral radius of the transition maps $\{\tau_\sigma\}_{\sigma \in \Sigma}$.

Bisimulation Pseudometric Between WFA

Definition (Difference automaton)

Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two weighted automata over the same finite alphabet Σ. Define their *difference automaton* as $\mathcal{A} = \mathcal{A}_1 - \mathcal{A}_2 = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ where $V = V_1 \oplus V_2$, $\alpha = \alpha_1 \oplus (-\alpha_2)$, $\beta = \beta_1 \oplus \beta_2$, and $\tau_\sigma = \tau_{1,\sigma} \oplus \tau_{2,\sigma}$ for all $\sigma \in \Sigma$.

Bisimulation Pseudometric Between WFA

Definition (Difference automaton)

Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two weighted automata over the same finite alphabet Σ. Define their *difference automaton* as $\mathcal{A} = \mathcal{A}_1 - \mathcal{A}_2 = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ where $V = V_1 \oplus V_2$, $\alpha = \alpha_1 \oplus (-\alpha_2)$, $\beta = \beta_1 \oplus \beta_2$, and $\tau_\sigma = \tau_{1,\sigma} \oplus \tau_{2,\sigma}$ for all $\sigma \in \Sigma$.

Definition

Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two weighted automata and let $\mathcal{A}$ be their difference automaton. For any $\gamma < 1/\rho(\mathcal{A})$ we define the γ-*bisimulation distance* between $\mathcal{A}_1$ and $\mathcal{A}_2$ as $d_\gamma(\mathcal{A}_1, \mathcal{A}_2) = s_{\mathcal{A},\gamma}(\alpha)$.

Bisimulation Pseudometric Between WFA

Proposition

Let $\mathcal{A}_1$ and $\mathcal{A}_2$ two weighted automata and $\gamma < 1/\max\{\rho(\mathcal{A}_1), \rho(\mathcal{A}_2)\}$. Then $d_\gamma(\mathcal{A}_1, \mathcal{A}_2)$ satisfies $d_\gamma(\mathcal{A}_1, \mathcal{A}_2) = 0$ if and only if $f_{\mathcal{A}_1} = f_{\mathcal{A}_2}$.

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.
- If we define a sequence of automata $(\mathcal{A}_i)$ that converges to an automaton $\mathcal{A}$, we want their distance to converge to zero.

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.
- If we define a sequence of automata $(\mathcal{A}_i)$ that converges to an automaton $\mathcal{A}$, we want their distance to converge to zero.

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.
- If we define a sequence of automata $(\mathcal{A}_i)$ that converges to an automaton $\mathcal{A}$, we want their distance to converge to zero. *(Parameter Continuity)*

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.
- If we define a sequence of automata $(\mathcal{A}_i)$ that converges to an automaton $\mathcal{A}$, we want their distance to converge to zero. *(Parameter Continuity)*
- An upper bound on the behavioural distance between two systems should imply an upper bound on the difference of their outputs as a function of the length of the input string.

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.
- If we define a sequence of automata $(\mathcal{A}_i)$ that converges to an automaton $\mathcal{A}$, we want their distance to converge to zero. *(Parameter Continuity)*
- An upper bound on the behavioural distance between two systems should imply an upper bound on the difference of their outputs as a function of the length of the input string.

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.
- If we define a sequence of automata $(\mathcal{A}_i)$ that converges to an automaton $\mathcal{A}$, we want their distance to converge to zero. *(Parameter Continuity)*
- An upper bound on the behavioural distance between two systems should imply an upper bound on the difference of their outputs as a function of the length of the input string. *(Input Continuity)*

Continuity Properties

- "Sanity check" for our bisimulation pseudometric.
- If we define a sequence of automata $(\mathcal{A}_i)$ that converges to an automaton $\mathcal{A}$, we want their distance to converge to zero. *(Parameter Continuity)*
- An upper bound on the behavioural distance between two systems should imply an upper bound on the difference of their outputs as a function of the length of the input string. *(Input Continuity)*
- Inspired by the continuity properties for labelled Markov chains presented by Jaeger et. al (2014).

Parameter Continuity

Definition

Let $(\mathcal{A}_i)_{i \in \mathbb{N}}$ be a sequence of WFA $\mathcal{A}_i = \langle \Sigma, V, \alpha_i, \beta_i, \{\tau_{i,\sigma}\}_{\sigma \in \Sigma} \rangle$ over the same alphabet Σ and normed vector space $(V, \|\cdot\|)$. We say that the sequence $(\mathcal{A}_i)$ *converges* to $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ if

- $\lim_{i \to \infty} \|\alpha_i - \alpha\| = 0$,
- $\lim_{i \to \infty} \|\beta_i - \beta\|_* = 0$,
- $\lim_{i \to \infty} \|\tau_{i,\sigma} - \tau_\sigma\|_{\mathrm{op}} = 0$ for all $\sigma \in \Sigma$.

Parameter Continuity

Definition

Let $(\mathcal{A}_i)_{i \in \mathbb{N}}$ be a sequence of WFA $\mathcal{A}_i = \langle \Sigma, V, \alpha_i, \beta_i, \{\tau_{i,\sigma}\}_{\sigma \in \Sigma} \rangle$ over the same alphabet Σ and normed vector space $(V, \|\cdot\|)$. We say that the sequence $(\mathcal{A}_i)$ *converges* to $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$ if

- $\lim_{i \to \infty} \|\alpha_i - \alpha\| = 0$,
- $\lim_{i \to \infty} \|\beta_i - \beta\|_* = 0$,
- $\lim_{i \to \infty} \|\tau_{i,\sigma} - \tau_\sigma\|_{\mathrm{op}} = 0$ for all $\sigma \in \Sigma$.

Definition

A pseudometric d between weighted automata is *parameter continuous* if for any sequence $(\mathcal{A}_i)_{i \in \mathbb{N}}$ converging to some weighted automaton $\mathcal{A}$, $\lim_{i \to \infty} d(\mathcal{A}, \mathcal{A}_i) = 0$.

Parameter Continuity

Theorem

The γ-bisimulation distance between weighted automata is parameter continuous for any sequence of weighted automata $(\mathcal{A}_i)_{i\in\mathbb{N}}$ converging to a weighted automaton $\mathcal{A}$ with $\gamma < 1/\rho(\mathcal{A})$.

Computing the Pseudometric

- Closed form expression for the seminorm:

$$s_{\mathcal{A},\gamma}(v) = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t |\beta(\tau_{x_{\leq t}}(v))| = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t |f_{\mathcal{A}_v}(x_{\leq t})| \ ,$$

and for the pseudometric:

$$d_\gamma(\mathcal{A}_1, \mathcal{A}_2) = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t \, |f_{\mathcal{A}_1}(x_{\leq t}) - f_{\mathcal{A}_2}(x_{\leq t})| \ .$$

Computing the Pseudometric

- Closed form expression for the seminorm:

$$s_{\mathcal{A},\gamma}(v) = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t |\beta(\tau_{x_{\leq t}}(v))| = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t |f_{\mathcal{A}_v}(x_{\leq t})| \ ,$$

  and for the pseudometric:

$$d_\gamma(\mathcal{A}_1, \mathcal{A}_2) = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t \left| f_{\mathcal{A}_1}(x_{\leq t}) - f_{\mathcal{A}_2}(x_{\leq t}) \right| \ .$$

- Supremum over all infinite strings and absolute value: looks hard to compute.

Undecidability Result

Theorem

The following problem is undecidable: given a weighted automaton $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$*, a discount factor* $\gamma < 1/\rho(\mathcal{A})$*, and a threshold* $\nu > 0$*, decide whether* $s_{\mathcal{A},\gamma}(\alpha) > \nu$*.*

Undecidability Result

Theorem

The following problem is undecidable: given a weighted automaton $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$, a discount factor $\gamma < 1/\rho(\mathcal{A})$, and a threshold $\nu > 0$, decide whether $s_{\mathcal{A},\gamma}(\alpha) > \nu$.

Proof idea: Reduction from computing the value function of unobservable MDPs (special cases of POMDPs) in a discounted infinite-horizon setting.

Spectral Learning Applications

- Despite the undecidability result, it is possible to bound the error of some algorithms in terms of the pseudometric to analyse their output.

Spectral Learning Applications

- Despite the undecidability result, it is possible to bound the error of some algorithms in terms of the pseudometric to analyse their output.
- The spectral learning algorithm from [Balle et al. 2014] PAC-learns WFA in terms of the pseudometric.

Spectral Learning Applications

- Despite the undecidability result, it is possible to bound the error of some algorithms in terms of the pseudometric to analyse their output.
- The spectral learning algorithm from [Balle et al. 2014] PAC-learns WFA in terms of the pseudometric.
  - With poly($1/\epsilon$) samples gives a hypothesis with error $< \epsilon$ measured with d_γ.

Spectral Learning Applications

- Despite the undecidability result, it is possible to bound the error of some algorithms in terms of the pseudometric to analyse their output.
- The spectral learning algorithm from [Balle et al. 2014] PAC-learns WFA in terms of the pseudometric.
  - With poly$(1/\epsilon)$ samples gives a hypothesis with error $< \epsilon$ measured with d_γ.
  - Previous results all use ℓ_1 distance on strings of bounded length, which is weaker.

Spectral Learning Applications

- Despite the undecidability result, it is possible to bound the error of some algorithms in terms of the pseudometric to analyse their output.
- The spectral learning algorithm from [Balle et al. 2014] PAC-learns WFA in terms of the pseudometric.
  - With poly($1/\epsilon$) samples gives a hypothesis with error $< \epsilon$ measured with d_γ.
  - Previous results all use ℓ_1 distance on strings of bounded length, which is weaker.
  - Proof idea: we combine continuity properties of pseudometric and continuity properties of joint spectral radius – this involves some delicate technical bounds.

Conclusion

- We constructed bisimulation pseudometrics between WFA and metrics between weighted languages.

Conclusion

- We constructed bisimulation pseudometrics between WFA and metrics between weighted languages.
- Satisfies parameter continuity and input continuity properties (under certain assumptions).

Conclusion

- We constructed bisimulation pseudometrics between WFA and metrics between weighted languages.
- Satisfies parameter continuity and input continuity properties (under certain assumptions).
- Applications to spectral learning.

Conclusion

- We constructed bisimulation pseudometrics between WFA and metrics between weighted languages.
- Satisfies parameter continuity and input continuity properties (under certain assumptions).
- Applications to spectral learning.
- Future work: develop an algorithm to approximately compute the bisimulation pseudometrics.

Conclusion

- We constructed bisimulation pseudometrics between WFA and metrics between weighted languages.
- Satisfies parameter continuity and input continuity properties (under certain assumptions).
- Applications to spectral learning.
- Future work: develop an algorithm to approximately compute the bisimulation pseudometrics.
  - Will most likely rely on the sum-of-squares programming approximation algorithm to compute the JSR of a set of matrices.

Thank you!

Appendix – Joint Spectral Radius

Definition

The *joint spectral radius* of a collection $M = \{\tau_i\}_{i \in I}$ of linear maps $\tau_i : V \to V$ on a normed vector space $(V, \|\cdot\|)$ is defined as

$$\rho(M) = \limsup_{t \to \infty} \left(\sup_{T \in I^t} \left\| \prod_{i \in T} \tau_i \right\| \right)^{1/t} = \lim_{t \to \infty} \left(\sup_{T \in I^t} \left\| \prod_{i \in T} \tau_i \right\| \right)^{1/t} .$$

Appendix – Joint Spectral Radius

Definition

The *joint spectral radius* of a collection $M = \{\tau_i\}_{i \in I}$ of linear maps $\tau_i : V \to V$ on a normed vector space $(V, \|\cdot\|)$ is defined as

$$\rho(M) = \limsup_{t\to\infty} \left(\sup_{T \in I^t} \left\| \prod_{i \in T} \tau_i \right\| \right)^{1/t} = \lim_{t\to\infty} \left(\sup_{T \in I^t} \left\| \prod_{i \in T} \tau_i \right\| \right)^{1/t} .$$

The joint spectral radius of $\mathcal{A}$, denoted as $\rho(\mathcal{A})$ is defined as $\rho(\{\tau_\sigma\}_{\sigma\in\Sigma})$ and can be rewritten as:

$$\rho(\mathcal{A}) = \lim_{t\to\infty} \left(\sup_{x \in \Sigma^t} \|\tau_x\| \right)^{1/t} .$$

Appendix – Metric for Banach's Fixed-Point Theorem

We define the following metric d on S:

$$d(s, s') = \sup_{\|v\| \leq 1} |s(v) - s'(v)| \ ,$$

where $\|\cdot\|$ is the following norm:

Theorem (Rota 1960)

Let $M = \{\tau_i\}_{i \in I}$ be a compact set of linear maps on V. For any $\eta > 0$ there exists a norm $\|\cdot\|$ on V that satisfies $\|\tau_i(v)\| \leq (\rho(M) + \eta) \|v\|$ for every $i \in I$ and every $v \in V$.

Appendix – Closed form

- Let $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$.
- $\gamma < 1/\rho(\mathcal{A})$.
- $F_{\mathcal{A},\gamma}(s)(v) = |\beta(v)| + \gamma \max_{\sigma \in \Sigma} s(\tau_\sigma(v))$.

Appendix – Closed form

- Let $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$.
- $\gamma < 1/\rho(\mathcal{A})$.
- $F_{\mathcal{A},\gamma}(s)(v) = |\beta(v)| + \gamma \max_{\sigma \in \Sigma} s(\tau_\sigma(v))$.

Theorem

$F_{\mathcal{A},\gamma}$ has a unique fixed point.

Appendix – Closed form

- Let $\mathcal{A} = \langle \Sigma, V, \alpha, \beta, \{\tau_\sigma\}_{\sigma \in \Sigma} \rangle$.
- $\gamma < 1/\rho(\mathcal{A})$.
- $F_{\mathcal{A},\gamma}(s)(v) = |\beta(v)| + \gamma \max_{\sigma \in \Sigma} s(\tau_\sigma(v))$.

Theorem

$F_{\mathcal{A},\gamma}$ has a unique fixed point.

Theorem

Let $s_{\mathcal{A},\gamma} \in \mathcal{S}$ be the fixed point of $F_{\mathcal{A},\gamma}$. Then for any $v \in V$ we have

$$s_{\mathcal{A},\gamma}(v) = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t |\beta(\tau_{x_{\leq t}}(v))| = \sup_{x \in \Sigma^\infty} \sum_{t=0}^{\infty} \gamma^t |f_{\mathcal{A}_v}(x_{\leq t})| \ .$$

Appendix – Reduction details

Theorem (Theorem 4.4 in Madani03)

The following problem is undecidable: given a UMDP $\mathcal{U}$ and a threshold ν decide whether there exists a sequence of actions $x \in \Sigma^\infty$ such that $V_{\mathcal{U}}(x) > \nu$.

Appendix – Reduction details

Theorem (Theorem 4.4 in Madani03)

The following problem is undecidable: given a UMDP $\mathcal{U}$ and a threshold ν decide whether there exists a sequence of actions $x \in \Sigma^\infty$ such that $V_{\mathcal{U}}(x) > \nu$.

The UMDP $\mathcal{U} = \langle \Sigma, Q, \alpha, \{\beta_\sigma\}_{\sigma\in\Sigma}, \{T_\sigma\}_{\sigma\in\Sigma} \rangle$ constructed in the proof always has:

Appendix – Reduction details

Theorem (Theorem 4.4 in Madani03)

The following problem is undecidable: given a UMDP $\mathcal{U}$ and a threshold ν decide whether there exists a sequence of actions $x \in \Sigma^\infty$ such that $V_{\mathcal{U}}(x) > \nu$.

The UMDP $\mathcal{U} = \langle \Sigma, Q, \alpha, \{\beta_\sigma\}_{\sigma \in \Sigma}, \{T_\sigma\}_{\sigma \in \Sigma} \rangle$ constructed in the proof always has:

1. action-independent rewards: $\beta_\sigma = \beta$ for all $\sigma \in \Sigma$,

Appendix – Reduction details

Theorem (Theorem 4.4 in Madani03)

The following problem is undecidable: given a UMDP $\mathcal{U}$ and a threshold ν decide whether there exists a sequence of actions $x \in \Sigma^\infty$ such that $V_{\mathcal{U}}(x) > \nu$.

The UMDP $\mathcal{U} = \langle \Sigma, Q, \alpha, \{\beta_\sigma\}_{\sigma\in\Sigma}, \{T_\sigma\}_{\sigma\in\Sigma}\rangle$ constructed in the proof always has:

1. action-independent rewards: $\beta_\sigma = \beta$ for all $\sigma \in \Sigma$,
2. non-negative rewards: $\beta_\sigma(q) \geq 0$ for all $q \in Q$ and $\sigma \in \Sigma$.

Appendix – Reduction details

Theorem (Theorem 4.4 in Madani03)

The following problem is undecidable: given a UMDP $\mathcal{U}$ and a threshold ν decide whether there exists a sequence of actions $x \in \Sigma^\infty$ such that $V_\mathcal{U}(x) > \nu$.

The UMDP $\mathcal{U} = \langle \Sigma, Q, \alpha, \{\beta_\sigma\}_{\sigma \in \Sigma}, \{T_\sigma\}_{\sigma \in \Sigma} \rangle$ constructed in the proof always has:

1. action-independent rewards: $\beta_\sigma = \beta$ for all $\sigma \in \Sigma$,
2. non-negative rewards: $\beta_\sigma(q) \geq 0$ for all $q \in Q$ and $\sigma \in \Sigma$.

Appendix – Reduction details

Theorem (Theorem 4.4 in Madani03)

The following problem is undecidable: given a UMDP $\mathcal{U}$ and a threshold ν decide whether there exists a sequence of actions $x \in \Sigma^\infty$ such that $V_\mathcal{U}(x) > \nu$.

The UMDP $\mathcal{U} = \langle \Sigma, Q, \alpha, \{\beta_\sigma\}_{\sigma\in\Sigma}, \{T_\sigma\}_{\sigma\in\Sigma}\rangle$ constructed in the proof always has:

1. action-independent rewards: $\beta_\sigma = \beta$ for all $\sigma \in \Sigma$,
2. non-negative rewards: $\beta_\sigma(q) \geq 0$ for all $q \in Q$ and $\sigma \in \Sigma$.

Corollary

The problem in the previous theorem remains undecidable when restricted to UMDP with non-negative action-independent rewards.