Spectral Learning from a Single Trajectory under Finite-State Policies (Supplementary Material)

Borja Balle   Odalric-Ambrym Maillard

A. Mixing Properties of Probabilistic Automata

Lemma 2 (η-mixing for PFA) *Let $\mathbb{A}$ be PFA and assume that it is (C,θ)-geometrically mixing in the sense that for some constants $C > 0, \theta \in (0,1)$ we have*

$$\forall t \in \mathbb{N}, \quad \mu_t^{\mathbb{A}} = \sup_{\alpha,\alpha'} \frac{\|\alpha A^t - \alpha' A^t\|_1}{\|\alpha - \alpha'\|_1} \leqslant C\theta^t \,,$$

where the supremum is over all probability vectors. Then we have $\eta_{\rho_{\mathbb{A}}} \leqslant C/(\theta(1-\theta))$.

Proof of Lemma 2:

We start by controlling the term η, defined by

$$\eta_{\rho_{\mathbb{A}}} = 1 + \max_{1<i<t} \sum_{j=i+1}^{t} \eta_{i,j} \,,$$

We proceed similarly to Lemma 7 of (Kontorovich & Weiss, 2014). By definition of the total variation norm $\|\cdot\|_{TV}$,

$$\eta_{i,j} = \frac{1}{2} \sup_{u\in\Sigma^{i-1},\sigma,\sigma'\in\Sigma} \sup_{Z\subseteq\Sigma^{t-j+1}} \left| \frac{\alpha^\top A_u A_\sigma A^{j-i-1} A_Z \beta}{\alpha^\top A_u A_\sigma \beta} - \frac{\alpha^\top A_u A_{\sigma'} A^{j-i-1} A_Z \beta}{\alpha^\top A_u A_{\sigma'} \beta} \right| ,$$

where $A_Z = \sum_{z\in Z} A_z$. At this point, it is convenient to introduce the vector $\alpha_{u,\sigma}^\top = \frac{\alpha^\top A_u A_\sigma}{\alpha^\top A_u A_\sigma \beta}$. Indeed, we then have the rewriting

$$\begin{aligned} \eta_{i,j} &= \frac{1}{2} \sup_{u\in\Sigma^{i-1},\sigma,\sigma'\in\Sigma} \sup_{Z\subseteq\Sigma^{t-j+1}} \left| (\alpha_{u,\sigma} - \alpha_{u,\sigma'})^\top A^{j-i-1} A_Z \beta \right| \\ &\leqslant \frac{1}{2} \sup_{u\in\Sigma^{i-1},\sigma,\sigma'\in\Sigma} \sup_{Z\subseteq\Sigma^{t-j+1}} \|(\alpha_{u,\sigma} - \alpha_{u,\sigma'})^\top A^{j-i-1}\|_1 \|A_Z\beta\|_\infty \end{aligned}$$

where we used a simple application of Hölder inequality. Since $\mathbb{A}$ is a PFA, we note that $\|A_Z\beta\|_\infty \leqslant 1$ because $\|\sum_{|z|=t-j+1} A_z\beta\|_\infty = 1$ and all the entries are non-negative. Also note that $\alpha_{u,\sigma}^\top \beta = \|\alpha_{u,\sigma}^\top\|_1 = 1$. Thus $\|\alpha_{u,\sigma} - \alpha_{u,\sigma'}\| \leqslant 2$. We deduce from these steps that

$$\eta_{i,j} \leqslant \sup_{\alpha,\alpha'} \frac{\|(\alpha - \alpha')^\top A^{j-i-1}\|_1}{\|\alpha - \alpha'\|_1} \,,$$

where the supremum is taken over all α, α' that are probability vectors. We note that the later quantity is precisely the definition of the coefficient $\mu_{j-i-1}^{\mathbb{A}}$. Assuming (C,θ)-geometrically mixing, that is $\mu_j^{\mathbb{A}} \leqslant C\theta^j$ for all j, this implies that

$$\eta_{i,j} \leqslant C\theta^{j-i-1} \,.$$

We then deduce that

$$\eta_{\rho_{\mathbb{A}}} \leqslant 1 + C \max_{1<i<t} \sum_{j=i+1}^{t} \theta^{j-i-1} \leqslant \frac{C}{\theta}(1 + \sum_{j=1}^{t-2} \theta^j) = \frac{C}{\theta} \frac{1-\theta^{t-1}}{1-\theta} \leqslant \frac{C}{\theta(1-\theta)} . \qquad \square$$

The following result provides a control of the $\eta_{\rho_{\mathbb{A}}}$ coefficients, and shows this can be made explicit in specific cases.

Corollary 1 *Let $\mathbb{A}$ be PFA with n states and assume that its matrix A has a spectral gap, that is $|\lambda_2(A)| < 1$. then there exists C such that $\eta_{\rho_{\mathbb{A}}} \leqslant \frac{C}{|\lambda_2(A)|(1-|\lambda_2(A)|)}$. When the corresponding chain is further aperiodic, irreducible and reversible, we further have $C \leqslant \sqrt{n}$.*

Proof of Corollary 1:

The first part of the result is folklore, and can be proven using some tedious steps involving the Jordan decomposition of the matrix see e.g. Fact 3 in (Rosenthal, 1995).

When the chain is irreducible, aperiodic and more importantly reversible, the spectral gap admits the following characterization, see Lemma 2.2 from (Kontoyiannis & Meyn, 2012):

$$\gamma_2(A) = \lambda_2(A) = \sup\left\{ \frac{\|A\nu\|_2}{\|\nu\|_2} : \nu \text{ s.t. } \|\nu\|_2 \neq 0, \nu^\top \mathbf{1} = 0 \right\}.$$

Thus, from $\lambda_2(A) < 1$ together with a change of norm from $\|\cdot\|_1$ to $\|\cdot\|_2$ and a standard argument (closely following that of Lemma 7), we obtain that

$$\mu_j^{\mathbb{A}} \leqslant C|\lambda_2|^j ,$$

where $C = \max_{x\in\mathbb{R}^n} \frac{||x||_1}{||x||_2} = \sqrt{n}$. $\square$

We end this section with a more technical lemma, that is useful to decompose terms in the proof of Theorem 3.

Lemma 6 (Mixing times of PFA) *Let $\mathbb{A} = \langle \alpha, \beta, \{A_\sigma\}\rangle$ be a PFA. Then, for any $s \geqslant s' \in \mathbb{N}$ it holds*

$$\|A^{s'} - \beta\alpha^\top A^s\|_\infty \leqslant 2\mu_{s'}^{\mathbb{A}} .$$

Proof of Lemma 6:

Let denote $\alpha_s^\top = \alpha^\top A^s$. We need to bound $\|A^{s'} - \beta\alpha_s^\top\|_\infty$. Recall that for any matrix M the $\|\cdot\|_\infty$-induced norm is given by $\|M\|_\infty = \max_i \sum_j |M(i,j)| = \max_i \|M(i,:)\|_1$. The ith row of $A^{s'} - \beta\alpha_s^\top$ is given by $e_i^\top A^{s'} - \alpha_s^\top$, where e_i is the ith column of the identity matrix. In particular, $e_i^\top A^{s'}$ is the distribution over states after starting in state i and running the chain for s' steps, and $\alpha_s^\top$ is the distribution over states starting from the distribution given by α and running the chain for s steps. The latter can also be rewritten as $\alpha_s^\top = \alpha^\top A^s = \alpha^\top A^{s-s'} A^{s'} = \alpha_{s-s'}^\top A^{s'}$, where $\alpha_{s-s'}^\top$ is again a distribution over states. Therefore we obtain the desired bound, since:

$$\begin{aligned}
\|A^{s'} - \beta\alpha_s^\top\|_\infty &= \max_{i\in[n]} \|e_i^\top A^{s'} - \alpha_{s-s'}^\top A^{s'}\|_1 \\
&\leqslant \sup_{\alpha_1,\alpha_2} \frac{\|\alpha_1^\top A^{s'} - \alpha_2^\top A^{s'}\|_1}{\|\alpha_1 - \alpha_2\|_1} \|e_i - \alpha_{s-s'}\|_1 \\
&\leqslant 2\mu_{s'}^{\mathbb{A}} .
\end{aligned}$$

$\square$

B. Geometry of Stochastic Weighted Automata

Lemma 3 (claim (i)) *For any $w \in \Sigma^\star$ we have $\|A_w \beta\|_\beta \leqslant 1$.*

Proof of Lemma 3 (claim (i)):

We shall use the cone monotonicity property of $\|\cdot\|_\beta$, which says that $0 \leqslant_\mathcal{K} u \leqslant_\mathcal{K} v$ implies $\|u\|_\beta \leqslant \|v\|_\beta$. First note that by construction of $\mathcal{K}$ we have $0 \leqslant_\mathcal{K} A_w \beta$. If we show that $A_w \beta \leqslant_\mathcal{K} \beta$ also holds, then cone monotonicity implies $\|A_w \beta\|_\beta \leqslant \|\beta\|_\beta = 1$.

To prove the claim note that because β is an eigenvector of A of eigenvalue 1 we have $\beta = A^t \beta = \sum_{|w|=t} A_w \beta$. Therefore, $\beta - A_w \beta = \sum_{|w'|=|w|, w' \neq w'} A_{w'} \beta$ which is a vector in $\mathcal{K}$ because convex cones are closed under non-negative linear combinations, and we conclude that $A_w \beta \leqslant_\mathcal{K} \beta$. □

Lemma 3 (claim (ii)) *For any $w \in \Sigma^\star$ we have $\|\alpha^\top A_w\|_{\beta,*} = \alpha^\top A_w \beta$.*

Proof of Lemma 3 (claim (ii)):

By unrolling the definitions of the dual norm and B_β we get

$$\|\alpha^\top A_w\|_{\beta,*} = \sup_{-\beta \leqslant_\mathcal{K} v \leqslant_\mathcal{K} \beta} \alpha^\top A_w v \ .$$

Now note that for any v such that $\beta - v \in \mathcal{K}$ we have

$$\alpha^\top A_w v = \alpha^\top A_w \beta - \alpha^\top A_w (\beta - v) \leqslant \alpha^\top A_w \beta \ ,$$

where we used that $\beta - v \in \mathcal{K}$ implies $A_w(\beta - v) \in \mathcal{K}$ implies $\alpha^\top A_w (\beta - v) \geqslant 0$. Since $-\beta \leqslant_\mathcal{K} \beta \leqslant_\mathcal{K} \beta$, the supremum in the definition of $\|\alpha^\top A_w\|_{\beta,*}$ is attained at $v = \beta$ and the result follows. □

C. Mixing Properties of Stochastic Weighted Automata

Lemma 4 (η-mixing for SWFA) *Let $\mathbb{A}$ be SWFA and assume that it is (C, θ)-geometrically mixing in the sense that for some $C \geqslant 0, \theta \in (0,1)$,*

$$\mu_t^\mathbb{A} = \sup_{\alpha_0, \alpha_1 : \alpha_0^\top \beta = \alpha_1^\top \beta = 1} \frac{\|\alpha_0^\top A^t - \alpha_1^\top A^t\|_{\beta,*}}{\|\alpha_0 - \alpha_1\|_{\beta,*}} \leqslant C \theta^t \ .$$

Then the η-mixing coefficient satisfies

$$\eta_{\rho_\mathbb{A}} \leqslant \frac{C}{\theta(1-\theta)} \ .$$

Proof of Lemma 4:

The proof follows the same initial steps as for Lemma 2. Introducing the vector $\alpha_{u,\sigma}^\top = \frac{\alpha^\top A_u A_\sigma}{\alpha^\top A_u A_\sigma \beta}$, we then have the rewriting

$$\begin{aligned}\eta_{i,j} &= \frac{1}{2}\sup_{u\in\Sigma^{i-1},\sigma,\sigma'\in\Sigma}\sup_{Z\subseteq\Sigma^{t-j+1}}\left|(\alpha_{u,\sigma}-\alpha_{u,\sigma'})^\top A^{j-i-1}A_Z\beta\right| \\ &\leqslant \frac{1}{2}\sup_{u\in\Sigma^{i-1},\sigma,\sigma'\in\Sigma}\sup_{Z\subseteq\Sigma^{t-j+1}}\|(\alpha_{u,\sigma}-\alpha_{u,\sigma'})^\top A^{j-i-1}\|_{\beta,*}\|A_Z\beta\|_\beta\end{aligned}$$

where we used a simple application of Hölder inequality and the norm induced by β. Since $\mathbb{A}$ is a SWFA, the same argument in the proof of Lemma 3 (i) can be used to show that $\|A_Z\beta\|_\beta \leqslant 1$ for any $Z \subseteq \Sigma^{t-j+1}$. On the other hand, from Lemma 3 (ii) we have $1 = \alpha_{u,\sigma}^\top\beta = \|\alpha_{u,\sigma}^\top\|_{\beta,*}$. Thus $\|\alpha_{u,\sigma}-\alpha_{u,\sigma'}\|_{\beta,*} \leqslant 2$. We deduce from these steps that

$$\eta_{i,j} \leqslant \sup_{\alpha,\alpha'}\frac{\|(\alpha-\alpha')^\top A^{j-i-1}\|_{\beta,*}}{\|\alpha-\alpha'\|_{\beta,*}}\,,$$

where the supremum is taken over all α,α' that satisfy $\alpha^\top\beta = 1$ We note that the later quantity is precisely the definition of the coefficient $\mu_{j-i-1}^{\mathbb{A}}$. We then conclude similarly to the proof of Lemma 2. □

Lemma 7 (Geometrical mixing of weighted automata) *Let $\mathbb{A} = \langle\alpha,\beta,\{A_\sigma\}\rangle$ be a stochastic WFA, $A = \sum_\sigma A_\sigma$, and*

$$\gamma_\beta(A) = \sup\left\{\frac{\|A\nu\|_{\beta,*}}{\|\nu\|_{\beta,*}} : \nu \text{ s.t. } \|\nu\|_{\beta,*}\neq 0, \nu^\top\beta = 0\right\}.$$

be its spectral gap with respect to β. It holds that

$$\mu_t^{\mathbb{A}} = \sup_{\alpha_0,\alpha_1:\alpha_0^\top\beta=\alpha_1^\top\beta=1}\frac{\|\alpha_0^\top A^t - \alpha_1^\top A^t\|_{\beta,*}}{\|\alpha_0-\alpha_1\|_{\beta,*}} \leqslant \gamma_\beta(A)^t\ .$$

Proof of Lemma 7:

To this end, note that if α_0,α_1 are are such that $\alpha_0^\top\beta = \alpha_1^\top\beta = 1$, then $v = \alpha_0-\alpha_1$ is such that $v^\top\beta = 0$. A crucial remark is that since A is a weighted automaton matrix, $\alpha_0^\top A\beta = \alpha_1^\top A\beta = 1$ and thus $w = A(\alpha_0-\alpha_1)$ also satisfies $w^\top\beta = 0$. Likewise, $(\alpha_0-\alpha_1)^\top A^t\beta = 0$ for all $t\in\mathbb{N}$.

A second remark is that if $\|A^s v\|_{\beta,*} = 0$ for some $s < t$, then $\|A^t v\|_{\beta,*} = 0$. Thus, we can restrict to v such that $\|A^s v\|_{\beta,*} \neq 0$ for all $s \leqslant t$. Then, it comes for such $v = \alpha_0-\alpha_1$,

$$\frac{\|A^t\nu\|_{\beta,*}}{\|\nu\|_{\beta,*}} = \frac{\|AA^{t-1}\nu\|_{\beta,*}}{\|A^{t-1}\nu\|_{\beta,*}}\cdots\frac{\|A\nu\|_{\beta,*}}{\|\nu\|_{\beta,*}} \leqslant \gamma_\beta(A)^t\,.$$

For the last inequality, we used the fact that since A is a weighted automaton matrix, and $v = \alpha_0-\alpha_1$, then $v^\top A^s\beta = 0$ for all s. This guarantees that indeed $\frac{\|AA^s\nu\|_{\beta,*}}{\|A^s\nu\|_{\beta,*}} \leqslant \gamma_\beta(A)$ for all s. □

Lemma 8 (Mixing times of SWA) *Let $\mathbb{A} = \langle\alpha,\beta,\{A_\sigma\}\rangle$ be a SWFA. Then, for all $s \geqslant s' \in \mathbb{N}$ it holds*

$$\|A^{s'} - \beta\alpha^\top A^s\|_\beta \leqslant 2\mu_{s'}^{\mathbb{A}}\,.$$

Proof of Lemma 8:

Let $\alpha_s^\top = \alpha^\top A^s$. To prove this we proceed as follows:

$$\begin{aligned}\|A^{s'} - \beta\alpha_s^\top\|_\beta &= \sup_{\|v\|_\beta \leqslant 1} \|(A^{s'} - \beta\alpha_s^\top)v\|_\beta \\ &= \sup_{\|v\|_\beta \leqslant 1} \sup_{\|u\|_{\beta,*} \leqslant 1} u^\top(A^{s'} - \beta\alpha_s^\top)v \\ &= \sup_{\|u\|_{\beta,*} \leqslant 1} \|u^\top(A^{s'} - \beta\alpha_s^\top)\|_{\beta,*} \\ &= \sup_{\|u\|_{\beta,*} \leqslant 1} \|u^\top(A^{s'} - \beta\alpha_{s-s'}^\top A^{s'})\|_{\beta,*} \ .\end{aligned}$$

Next we note that for any u such that $\|u\|_{\beta,*} \leqslant 1$ we have $|u^\top\beta| \leqslant 1$, so:

$$\begin{aligned}\|u^\top\beta\alpha_t^\top\|_{\beta,*} &= |u^\top\beta|\|\alpha_t^\top\|_{\beta,*} \\ &\leqslant \|\alpha_t^\top\|_{\beta,*} \ .\end{aligned}$$

Furthermore, the same argument we used to show that $\|\alpha^\top A_x\|_{\beta,*} = \alpha^\top A_x\beta$ implies that $\|\alpha_t^\top\|_{\beta,*} = \|\alpha^\top A^t\|_{\beta,*} = \alpha^\top A^t\beta = 1$. Therefore, we see that $\|u\|_{\beta,*} \leqslant 1$ implies $\|u^\top\beta\alpha_t^\top\|_{\beta,*} \leqslant 1$, and we get the inequality

$$\begin{aligned}\|A^{s'} - \beta\alpha_s^\top\|_\beta &\leqslant \sup_{\|u_1\|_{\beta,*} \leqslant 1} \sup_{\|u_2\|_{\beta,*} \leqslant 1} \|u_1^\top A^{s'} - u_2^\top A^{s'}\|_{\beta,*} \\ &\leqslant \mu_{s'}^{\mathbb{A}} \sup_{\|u_1\|_{\beta,*} \leqslant 1} \sup_{\|u_2\|_{\beta,*} \leqslant 1} \|u_1 - u_2\|_{\beta,*} \\ &\leqslant 2\mu_{s',\beta}^{\mathbb{A}} \ .\end{aligned}$$

□

D. Single-Trajectory Concentration Inequalities for Probabilistic Automata

Theorem 2 (Single-trajectory, entry-wise concentration) *Let $\mathbb{A}$ be a PFA that is (C,θ)-geometrically mixing, and $\xi \sim \rho_{\mathbb{A}} \in \mathcal{P}(\Sigma^\omega)$ a trajectory of observations. Then for any $u \in \mathcal{U}, v \in \mathcal{V}$ and $\delta \in (0,1)$ it holds*

$$\mathbb{P}\bigg(\widehat{H}_{t,\xi}^{\mathcal{U},\mathcal{V}}(u,v) - \bar{H}_t^{\mathcal{U},\mathcal{V}})u,v) > \frac{|uv|C}{\theta(1-\theta)}\sqrt{\Big(1 + \frac{|uv|-1}{t}\Big)\frac{\ln(1/\delta)}{2t}}\bigg) \leqslant \delta \, .$$

Proof of Theorem 2:

We control $\eta_{\rho_{\mathbb{A}}}$ by a direct application of Lemma 2.

Control of $\|g\|_{Lip}$: Let us fix $u \in \mathcal{U}, v \in \mathcal{V}$ define $g(\xi) = t\widehat{H}_{t,\xi}^{\mathcal{U},\mathcal{V}}(u,v)$. We first control the regularity of f.

To this end, let ξ' be a trajectory $\xi' = x_1, \dots, x_{k-1}, x_k', x_{k+1}, \dots, x_\ell$ that only differs by one element from ξ, say at position k. Then, we get for any u, v

$$\begin{aligned}g(\xi) - g(\xi') &= \sum_{s=1}^{t}(\mathbb{I}\{x_s \dots x_{s+|uv|-1} = uv\} - \mathbb{I}\{x_s \dots x_k' \dots x_{s+|uv|-1} = uv\}) \\ &\leqslant |\{s \in [1,t] : k \in [s : s+|uv|-1]\}| \, .\end{aligned}$$

Now, in order to bound $|\{s \in [1,t] : k \in [s : s + |uv| - 1]\}|$ note that $k \in [s : s + |uv| - 1]$ if and only if $s \leqslant k \leqslant s + |uv| - 1$. From the first inequality we see that $s \leqslant k$, and from the second one $s \geqslant k - |uv| + 1$. Combined with the restrictions on s, this means that

$$|\{s \in [1,t] : k \in [s : s + |uv| - 1]\}| = |[\max\{1, k - |uv| + 1\}, \min\{k, t\}]| \leqslant |uv| \,,$$

which show that $\|g\|_{Lip} \leqslant |uv|$.

Combining the two quantities Combining these two results, and noting that $t + |uv| - 1$ symbols appears in $g(\xi)$, we deduce that $\forall \varepsilon > 0$,

$$\mathbb{P}\bigg(t(\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi}(u,v) - \bar{H}^{\mathcal{U},\mathcal{V}}_{t}(u,v)) > |uv|(t + |uv| - 1)\varepsilon \bigg) \leqslant \exp\bigg(- \frac{2(t + |uv| - 1)\theta^2(1-\theta)^2\varepsilon^2}{C^2} \bigg) ,$$

or equivalently, for all $\delta \in (0,1)$,

$$\mathbb{P}\bigg(\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi}(u,v) - \bar{H}^{\mathcal{U},\mathcal{V}}_{t}(u,v) > \frac{\sqrt{t + |uv| - 1}|uv|}{t} \frac{C}{\theta(1-\theta)} \sqrt{\frac{\ln(1/\delta)}{2t}} \bigg) \leqslant \delta \,. \qquad \square$$

The proof of following result is more challenging.

Theorem 3 (Single-trajectory, matrix-wise) *Let $\rho_{\mathbb{A}} \in \mathcal{P}(\Sigma^\omega)$ be as in Theorem 2 and define the probability mass $m^{\mathcal{U},\mathcal{V}} = \sum_{u\in\mathcal{U},v\in\mathcal{V}} \bar{f}_t(uv)$. Then, for all $\delta \in (0,1)$,*

$$\mathbb{P}\Bigg(\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t} - \bar{H}^{\mathcal{U},\mathcal{V}}_{t}\|_2 \geqslant \bigg(\sqrt{L} + \sqrt{\frac{2C}{1-\theta}}\bigg)\sqrt{\frac{2m^{\mathcal{U},\mathcal{V}}}{t}} + \frac{2LC}{\theta(1-\theta)}\sqrt{\bigg(1 + \frac{L-1}{t}\bigg)\frac{\min\{|\mathcal{U}||\mathcal{V}|, 2n_{\mathcal{U}}n_{\mathcal{V}}\}\ln(1/\delta)}{2t}} \Bigg) \leqslant \delta \,.$$

Proof of Theorem 3:

Let us introduce the function $g(\xi) = \|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t} - \bar{H}^{\mathcal{U},\mathcal{V}}_{t}\|_2$. We first control $\|g\|_{Lip}$ then $\mathbb{E}[g(\xi)]$, before applying Theorem 1.

Step 1: Control of $\|g\|_{Lip}$. In this step, we show that

$$\|g\|_{Lip} \leqslant \frac{L}{t}\sqrt{\min\{|\mathcal{U}||\mathcal{V}|, 2n_{\mathcal{U}}n_{\mathcal{V}}\}}$$

where $L = \max_{u\in\mathcal{U},v\in\mathcal{V}} |uv|$ denote the maximal length of words in $U \cdot V$ and $n_{\mathcal{U}} = |\ell \in [0,L] : |U_\ell| > 0|$, $n_{\mathcal{V}} = |\ell \in [0,L] : |V_\ell| > 0|$, denote the number lengths such that the set $U_\ell = \{u \in \mathcal{U} : |u| = \ell\}$ (respectively $V_\ell = \{v \in \mathcal{V} : |v| = \ell\}$) is non empty. Note that the second term in the min can be exponentially smaller than the first. For example, taking $U = V = \Sigma^{\leqslant L/2}$ we have $|U||V| = \Theta(|\Sigma|^L)$ while $n_{\mathcal{U}}n_{\mathcal{V}} = \Theta(L^2)$.

Step 1.1. Let $\xi' \sim p$ be a trajectory $\xi' = x_1, \ldots, x_{k-1}, x'_k, x_{k+1}, \ldots, x_\ell$ that only differs by one element from ξ, say at position k. We note that

$$\begin{aligned}
\Big| \|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \bar{H}^{\mathcal{U},\mathcal{V}}_{t}\|_2 - \|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi'} - \bar{H}^{\mathcal{U},\mathcal{V}}_{t}\|_2 \Big| &\leqslant \|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi'}\|_2 \\
= \sup_{q \in \mathbb{R}^V} \frac{1}{t\|q\|_2} \sqrt{\sum_{u \in \mathcal{U}} \Big(\sum_{v \in \mathcal{V}} \sum_{s=1}^{t} \left(\mathbb{I}\{x_s \dots x_{s+|uv|-1} = uv\} - \mathbb{I}\{x_s \dots x'_k \dots x_{s+|uv|-1} = uv\} \right) q_v \Big)^2} \\
\leqslant \sup_{q \in \mathbb{R}^V} \frac{1}{t\|q\|_2} \sqrt{\sum_{u \in \mathcal{U}} \Big(\sum_{v \in \mathcal{V}} |uv| q_v \Big)^2} \leqslant \sup_{q \in \mathbb{R}^V} \frac{1}{t\|q\|_2} \sqrt{\sum_{u \in \mathcal{U}} \sum_{v \in \mathcal{V}} |uv|^2} \sqrt{\sum_{v \in \mathcal{V}} q_v^2} \,.
\end{aligned}$$

Let $L = \max_{u \in \mathcal{U}, v \in \mathcal{V}} |uv|$. A simple bound is then $\|g\|_{Lip} \leqslant \frac{L}{t}\sqrt{|\mathcal{U}||\mathcal{V}|}$, which is essentially optimal if all words uv, $u \in \mathcal{U}, v \in \mathcal{V}$ have same length.

Step 1.2. A more refined bound may be helpful in case many words have length $|uv|$ much smaller than L. To his end, let us write $\widehat{H}^{\mathcal{U},\mathcal{V}}_{t} = \frac{1}{t}\sum_{s=1}^{t} M_s$ with $M_s(u,v) = b_{s,uv} = \mathbb{I}\{x_s \dots x_{s+|uv|-1} = uv\}$. Similarly, let $H^{\mathcal{U},\mathcal{V}}_{t,\xi'} = \frac{1}{t}\sum_{s=1}^{t} M'_s$ with the obvious definition. Now, by the same argument we used to bound $\|g\|_{Lip}$ in the entry-wise case, we have

$$t\left(\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi'}\right) = \sum_{s=k-L+1}^{k} M_s - M'_s = \sum_{s=k-L+1}^{k} \Delta_s \,,$$

since for $s < k - L + 1$ or $s > k$ we must have $M_s = M'_s$.

Now let us partition the sets U and V as disjoint unions of sets with strings of the same length. That is, we write $U = \cup_{\ell=0}^{L} U_\ell$ with $U_\ell = U \cap \Sigma^\ell$, and $V = \cup_{\ell=0}^{L} V_\ell$ with analogous definitions. This allows us to write $M_s \in \{0,1\}^{U \times V}$ as a block matrix $M_s = (M_s^{i,j})_{0 \leqslant i,j \leqslant L}$ with $M_s^{i,j} \in \{0,1\}^{U_i \times V_j}$.

For simplicity of notation, in the sequel we are assuming that $U_\ell, V_\ell \neq \emptyset$ for all $0 \leqslant \ell \leqslant L$, but the argument remains the same after we remove the empty sets of rows and columns. Note that by definition we have $M_s^{i,j}(u,v) = \mathbb{I}\{x_s \dots x_{s+i+j-1} = uv\}$ for any $u \in \mathcal{U}_i$ and $v \in \mathcal{V}_j$. This implies that each of the block matrices $M_s^{i,j}$ contains *at most one non-zero entry*.

If we make analogous definitions and write $M'_s = ({M'}_s^{i,j})_{0 \leqslant i,j \leqslant L}$, then we obtain a block decomposition for $\Delta_s = M_s - M'_s = (\Delta_s^{i,j})_{0 \leqslant i,j \leqslant L}$ where each block is either:

1. zero,
2. a $\{0,1\}$-matrix with a single 1,
3. a $\{0,-1\}$-matrix with a single -1,
4. a $\{0,1,-1\}$-matrix with a single 1 and a single -1.

In any of these cases one can see that the bound $\|\Delta_s^{i,j}\|_2 \leqslant \|\Delta_s^{i,j}\|_F \leqslant \sqrt{2}$ is always satisfied. Therefore, we have $\|\Delta_s\|_2^2 \leqslant \|\Delta_s\|_F^2 = \sum_{i,j} \|\Delta_s^{i,j}\|_F^2 \leqslant 2 n_{\mathcal{U}} n_{\mathcal{V}}$, where

$$\begin{aligned}
n_{\mathcal{U}} &= |\ell \in [0,L] : |U_\ell| > 0| \,, \\
n_{\mathcal{V}} &= |\ell \in [0,L] : |V_\ell| > 0| \,.
\end{aligned}$$

By plugging these estimates into $\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi'}$ we finally get

$$\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi'}\|_2 \leqslant \frac{L\sqrt{2 n_{\mathcal{U}} n_{\mathcal{V}}}}{t} \,.$$

Therefore we obtain the bound $\|g\|_{Lip} \leqslant (L\sqrt{2 n_{\mathcal{U}} n_{\mathcal{V}}})/t$.

Control of $\mathbb{E}[g(\xi)]$. We now want to control the following quantity $\mathbb{E}[\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2]$. More precisely, we show in this step that

$$\mathbb{E}[\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2]^2 \quad \leqslant \quad \frac{O\left(L^2 + \frac{1}{1-\theta}\right)\left(\sum_{u\in\mathcal{U},v\in\mathcal{V}} \bar{f}_t(uv)\right)}{t} \ ,$$

where $L = \max_{w\in\mathcal{U}\cdot V} |w|$, and $\bar{f}_t(w) = \frac{1}{t}\sum_{s=1}^t f_s(w)$, where $f_s(w) = \mathbb{P}[\xi \in \Sigma^{s-1}w\Sigma^\omega]$.

Step 2.1. Let $q \in \mathbb{R}^V$ be a unit vector ($\|q\|_2 = 1$.) Then, by Jensen's inequality, the norm of $\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}$ is controlled by its Frobenius norm

$$\begin{aligned}
\mathbb{E}[\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2]^2 \quad &\leqslant \quad \mathbb{E}[\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2^2] \\
&\leqslant \quad \mathbb{E}[\sum_{u\in\mathcal{U}}\Big(\sum_{v\in\mathcal{V}}(\widehat{H}_t^{\mathcal{U},\mathcal{V}}(u,v) - \bar{H}_t^{\mathcal{U},\mathcal{V}}(u,v))q_v\Big)^2] \\
&\leqslant \quad \sum_{u\in\mathcal{U}}\sum_{v\in\mathcal{V}}\mathbb{E}[(\widehat{H}_t^{\mathcal{U},\mathcal{V}}(u,v) - \bar{H}_t^{\mathcal{U},\mathcal{V}}(u,v))^2] \\
&= \quad \sum_{w\in\mathcal{U}\cdot V}|w|_{\mathcal{U},\mathcal{V}}\mathbb{E}[(\widehat{f}_t(w) - \bar{f}_t(w))^2] \ ,
\end{aligned}$$

where $U\cdot V$ is the set of all words of the form $u\cdot v$ with $u\in\mathcal{U}$ and $v\in\mathcal{V}$; $|w|_{\mathcal{U},\mathcal{V}} = |(u,v)\in\mathcal{U}\times V : u\cdot v = w|$, and $\widehat{f}_t(w) = \frac{1}{t}\sum_{s=1}^t b_{s,w}$ with the notation defined above. We also use $\bar{f}_t(w) = \mathbb{E}[\widehat{f}_t(w)] = \frac{1}{t}\sum_{s=1}^t f_s(w)$, where $f_s(w) = \mathbb{P}[\xi\in\Sigma^{s-1}w\Sigma^\omega]$. This implies that we have a sum of variances, and each of them can be written as

$$\mathbb{E}[(\widehat{f}_t(w) - \bar{f}_t(w))^2] = \mathbb{E}[\widehat{f}_t(w)^2] - \bar{f}_t(w)^2 \ .$$

An important first observation is that we can write $f_s(w) = \alpha^\top A^{s-1}A_w\beta$. Furthermore, it follows from A being a probabilistic automaton that $\sum_{|w|=l} f_s(w) = 1$ for all s and l. This suggests that we group the terms in the sum over $W = U\cdot V$ by length, so we write $W_l = W\cap\Sigma^l$ and define $L_l = \max_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}$ the maximum number of ways to write a string of length l in W as a product of a prefix in U and a suffix in V. Note that we always have $L_l \leqslant l+1$. Henceforth, we want to control the following terms for all possible values of l:

$$\sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}\Big(\mathbb{E}[\widehat{f}_t(w)^2] - \bar{f}_t(w)^2\Big) = \frac{1}{t^2}\sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}\left[\mathbb{E}\left[\left(\sum_{s=1}^t b_{s,w}\right)^2\right] - \left(\sum_{s=1}^t f_s(w)\right)^2\right].$$

Step 2.2. Let us focus on each of the quadratic terms. On the one hand, it holds

$$\left(\sum_{s=1}^t f_s(w)\right)^2 \quad = \quad \sum_{s=1}^t f_s(w)^2 + 2\sum_{1\leqslant s<s'\leqslant t} f_s(w)f_{s'}(w) \, ,$$

while other on the other hand, we get

$$\mathbb{E}\left[\left(\sum_{s=1}^t b_{s,w}\right)^2\right] \quad = \quad \sum_{s=1}^t \mathbb{E}[b_{s,w}^2] + 2\sum_{1\leqslant s<s'\leqslant t}\mathbb{E}[b_{s,w}b_{s',w}] \, .$$

Hence this enables to derive the following bound

$$\begin{aligned}
\mathbb{E}[\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2]^2 \quad \leqslant \quad & \frac{1}{t^2}\sum_{l=0}^\infty\sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}\Bigg[\sum_{s=1}^t(1-f_s(w))f_s(w) \\
& +2\sum_{1\leqslant s<s'\leqslant t}\Big(\mathbb{E}[b_{s,w}b_{s',w}] - f_s(w)f_{s'}(w)\Big)\Bigg] \, . \qquad (6)
\end{aligned}$$

Step 2.3. In order to control the first term in (6), we remark that

$$\begin{aligned}\sum_{l=0}^{\infty}\sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}\sum_{s=1}^{t}(1-f_s(w))f_s(w) &= \sum_{u\in\mathcal{U},v\in\mathcal{V}}\sum_{s=1}^{t}(1-f_s(uv))f_s(uv)\\ &\leqslant \sum_{u\in\mathcal{U},v\in\mathcal{V}}\sum_{s=1}^{t}f_s(uv)\\ &= t\sum_{u\in\mathcal{U},v\in\mathcal{V}}\bar{f}_t(uv)\,. \end{aligned}\tag{7}$$

Step 2.4. We thus focus on controlling the remaining "cross"-term in (6) and to this end we study, for $w \in W_l$, the quantity

$$\mathbb{E}[b_{s,w}b_{s',w}] - f_s(w)f_{s'}(w) = \mathbb{P}[\xi\in\Sigma^{s-1}\Sigma_w^{s'-s}\Sigma^{\omega}] - (\alpha^\top A^{s-1}A_w\beta)(\alpha^\top A^{s'-1}A_w\beta)\,,$$

where we introduced for convenience the set $\Sigma_w^{s'-s} = w\Sigma^{s'-s}\cap\Sigma^{s'-s}w$. Introducing as well the vectors $\alpha_{s-1}^\top = \alpha^\top A^{s-1}$, $\alpha_{s'-1}^\top = \alpha^\top A^{s'-1}$ and the transition matrix $A_w^{s'-s} = \sum_{x\in\Sigma_w^{s'-s}} A_x$ corresponding to the "event" $\Sigma_w^{s'-s}$, it comes

$$\mathbb{E}[b_{s,w}b_{s',w}] - f_s(w)f_{s'}(w) = \alpha_{s-1}^\top\Big(A_w^{s'-s} - A_w\beta\alpha_{s'-1}^\top A_w\Big)\beta\,.$$

We now discuss two cases. First the case when $s'-s\geqslant l$, then the case when $s'-s<l$.

Note that if $s'-s\geqslant |w| = l$, then $\Sigma_w^{s'-s}$ simplifies to $\Sigma_w^{s'-s} = w\Sigma^{s'-s-l}w$ and thus $A_w^{s'-s} = A_wA^{s'-s-l}A_w$. For such words, we thus obtain

$$\begin{aligned}\alpha_{s-1}^\top\Big(A_w^{s'-s} - A_w\beta\alpha_{s'-1}^\top A_w\Big)\beta &= \alpha_{s-1}^\top A_w\Big(A^{s'-s-l}-\beta\alpha_{s'-1}^\top\Big)A_w\beta\\ &\leqslant \|\alpha_{s-1}^\top A_w\|_1\|A^{s'-s-l}-\beta\alpha_{s'-1}^\top\|_\infty\|A_w\beta\|_\infty\,.\end{aligned}$$

Moreover, from Lemma 6, it holds $\|A^{s'-s-l}-\beta\alpha_{s'-1}^\top\|_\infty \leqslant 2\mu^{\mathbb{A}}_{s'-s-l}$. Also, it holds that $\|A_w\beta\|_\infty\leqslant 1$. Finally, since $\alpha_{s-1}^\top A_w$ is a sub-distribution over states, we have

$$\begin{aligned}\sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}\|\alpha_{s-1}^\top A_w\|_1 &= \sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}\alpha_{s-1}^\top A_w\beta\\ &= \sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}f_s(w) = \sum_{u\in\mathcal{U},v\in\mathcal{V}:uv\in W_l}f_s(uv)\ .\end{aligned}$$

Now, on the other hand if $s'-s<l$, using the fact that $\Sigma_w^{s'-s}\subset w\Sigma^{s'-s}$, then

$$\begin{aligned}\alpha_{s-1}^\top\Big(A_w^{s'-s} - A_w\beta\alpha_{s'-1}^\top A_w\Big)\beta &\leqslant \alpha_{s-1}^\top A_w\Big(A^{s'-s}-\beta\alpha_{s'-1}^\top A_w\Big)\beta\\ &= f_s(w)(1-f_{s'}(w))\leqslant f_s(w)\ .\end{aligned}$$

So in this case we again see that $\sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}f_s(w) = \sum_{u\in\mathcal{U},v\in\mathcal{V}:uv\in W_l}f_s(uv)$.

Step 2.5. Therefore, combining the above steps, so far we have seen that for a fixed $l\geqslant 0$, the sum $\sum_{w\in W_l}|w|_{\mathcal{U},\mathcal{V}}\sum_{1\leqslant s<s'\leqslant t}(\mathbb{E}[b_{s,w}b_{s',w}]-f_s(w)f_{s'}(w))$ is upper bounded by:

$$\begin{aligned}&\sum_{1\leqslant s<s'\leqslant t}\ \sum_{u\in\mathcal{U},v\in\mathcal{V}:|uv|=l} f_s(uv)(2\mu^{\mathbb{A}}_{s'-s-l}\mathbb{I}\{s'-s\geqslant l\}+\mathbb{I}\{s'-s<l\})\\ =& \sum_{u\in\mathcal{U},v\in\mathcal{V}:|uv|=l}\ \sum_{s=1}^{t-1}f_s(uv)\Bigg[\sum_{s'=s+1}^{t}2\mu^{\mathbb{A}}_{s'-s-l}\mathbb{I}\{s'-s\geqslant l\}+\mathbb{I}\{s'-s<l\}\Bigg]\ .\end{aligned}$$

Now note that $\sum_{s'=s+1}^{t} \mathbb{I}\{s'-s<l\} = \min\{l-1, t-s\} \leqslant l-1$. Furthermore, using that $\mu_t^{\mathbb{A}} \leqslant C\theta^t$ we get

$$\begin{aligned}\sum_{s'=s+1}^{t} \mu^{\mathbb{A}}_{s'-s-l}\mathbb{I}\{s'-s \geqslant l\} &= \mathbb{I}\{t \geqslant s+l\} \sum_{k=0}^{t-s-l} \mu_k^{\mathbb{A}} \\ &\leqslant C\mathbb{I}\{t \geqslant s+l\} \frac{1-\theta^{t-s-l+1}}{1-\theta} \leqslant \frac{C}{1-\theta} \ .\end{aligned}$$

In conclusion, we get

$$\begin{aligned}&\sum_{u\in\mathcal{U},v\in\mathcal{V}:|uv|=l} \sum_{s=1}^{t-1} f_s(uv)\Big[\sum_{s'=s+1}^{t} 2\mu^{\mathbb{A}}_{s'-s-l}\mathbb{I}\{s'-s \geqslant l\} + \mathbb{I}\{s'-s<l\}\Big] \\ \leqslant\ & \Big(l-1+\frac{2C}{1-\theta}\Big) \sum_{u\in\mathcal{U},v\in\mathcal{V}:|uv|=l} \sum_{s=1}^{t-1} f_s(uv) \\ \leqslant\ & t\Big(l-1+\frac{2C}{1-\theta}\Big) \sum_{u\in\mathcal{U},v\in\mathcal{V}:|uv|=l} \bar{f}_t(uv) \ .\end{aligned}$$

Finally, putting all the pieces together and introducing $L = \max_{w\in\mathcal{U}\cdot V} |w|$, we get from equations (6), (7), (8),

$$\begin{aligned}\mathbb{E}[\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2]^2 &\leqslant \frac{\sum_{u\in\mathcal{U},v\in\mathcal{V}} \bar{f}_t(uv)}{t} + \frac{2}{t}\sum_{l=0}^{\infty} \sum_{u\in\mathcal{U},v\in\mathcal{V}:|uv|=l} \bar{f}_t(uv)(l-1+\frac{2C}{1-\theta}) \\ &\leqslant \Big[2L-1+\frac{4C}{1-\theta}\Big]\frac{\sum_{u\in\mathcal{U},v\in\mathcal{V}} \bar{f}_t(uv)}{t} \ .\end{aligned}$$

Step 3. Application of Theorem 1. It remains to apply Theorem 1 with

$$\begin{aligned}\|g\|_{Lip} &\leqslant \frac{L}{t}\sqrt{\min\{|\mathcal{U}||\mathcal{V}|, 2n_{\mathcal{U}}n_{\mathcal{V}}\}}\,, \\ \mathbb{E}[\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2] &\leqslant \Big(\sqrt{L}+\sqrt{\frac{2C}{1-\theta}}\Big)\sqrt{\frac{2\sum_{u\in\mathcal{U},v\in\mathcal{V}} \bar{f}_t(uv)}{t}} \ ,\end{aligned}$$

for some constant C. After some rewriting, it comes

$$\begin{aligned}\mathbb{P}\Bigg(\|\widehat{H}_t^{\mathcal{U},\mathcal{V}} - \bar{H}_t^{\mathcal{U},\mathcal{V}}\|_2 > \Big(\sqrt{L}+\sqrt{\frac{2C}{1-\theta}}\Big)\sqrt{\frac{2\sum_{u\in\mathcal{U},v\in\mathcal{V}} \bar{f}_t(uv)}{t}} \\ + \frac{LC}{(1-\theta)}\sqrt{\Big(1+\frac{L-1}{t}\Big)\frac{\min\{|\mathcal{U}||\mathcal{V}|, 2n_{\mathcal{U}}n_{\mathcal{V}}\}\ln(1/\delta)}{2t}}\Bigg) \leqslant \delta\,.\end{aligned}$$

□

E. Single-Trajectory Hankel Concentration Inequalities with Finite-State Control

Lemma 5 *The Hankel matrix $\widehat{H} = \widehat{H}_{t,\xi}^{\mathcal{U},\mathcal{V}}$ computed in Algorithm 3 satisfies $\mathbb{E}[\widehat{H}_{t,\xi}^{\mathcal{U},\mathcal{V}}] = \tilde{H}_t^{\mathcal{U},\mathcal{V}}$, where $\tilde{H}_t^{\mathcal{U},\mathcal{V}}$ is a block of the Hankel matrix corresponding to the stochastic WFA $\tilde{\mathbb{A}}_t = \langle \tilde{\alpha}_t, \beta, \{A_\sigma\}\rangle$ where we introduced the modified vector $\tilde{\alpha}_t = (1/t)\sum_{s=0}^{t-1} \alpha^\top (A/\kappa)^s$. We denote by $\tilde{f}_t$ the function computed by $\tilde{\mathbb{A}}_t$.*

Proof of Lemma 5:

For any $t \geqslant 0$ and $w \in \Sigma^\star$ let us define the function $\varphi_{s,w} : \Sigma^\omega \to \mathbb{R}$ given by

$$\varphi_{s,w}(x) = \frac{\mathbb{I}\{o_{s+1}a_{s+1}\cdots o_{s+|uv|}a_{s+|uv|} = w\}}{\kappa^s \pi(a_1 \cdots a_{s+|w|}|o_1 \cdots o_{s+|w|})} \ ,$$

where $x = (o_1, a_1)(o_2, a_2)\cdots$. Thus, the entries of the Hankel matrix computed in Algorithm 3 can be written as $\widehat{H}(u,v) = (1/t)\sum_{s=0}^{t-1}\varphi_{s,uv}(\xi)$. Now note that the expectation $\mathbb{E}[\varphi_{s,w}]$ with respect to a trajectory $\xi \sim \rho_{\mathbb{B}}$ can be written as

$$\sum_{w' \in \Sigma^s} \frac{\mathbb{P}[\xi \in w'w\Sigma^\omega]}{\kappa^s \pi(w'^{\mathcal{A}} w^{\mathcal{A}} | w'^{\mathcal{O}} w^{\mathcal{O}})} = \sum_{w' \in \Sigma^s} \frac{f_{\mathbb{B}}(w'w)}{\kappa^s f_{\mathbb{A}_\pi}(w'w)}$$
$$= \sum_{w' \in \Sigma^s} \frac{f_{\mathbb{A}}(w'w)}{\kappa^s} = \frac{\alpha^\top A^s A_w \beta}{\kappa^s} \ .$$

Therefore, the Hankel matrix $\widehat{H} = \widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi}$ computed in Algorithm 3 satisfies $\mathbb{E}[\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi}] = \tilde{H}^{\mathcal{U},\mathcal{V}}_t$, where $\tilde{H}^{\mathcal{U},\mathcal{V}}_t$ is a block of the Hankel matrix corresponding to the stochastic WFA $\tilde{\mathbb{A}}_t = \langle \tilde{\alpha}_t, \beta, \{A_\sigma\}\rangle$ with modified vector $\tilde{\alpha}_t = (1/t)\sum_{s=0}^{t-1}\alpha^\top (A/\kappa)^s$. We denote by $\tilde{f}_t$ the function computed by $\tilde{\mathbb{A}}_t$. $\square$

Theorem 6 (Controlled case, single-trajectory, matrix-wise) *Let $\mathbb{A} = \langle \alpha, \beta, \{A_\sigma\}\rangle$ be a stochastic environment and π a stochastic policy induced by a probabilistic automaton $\mathbb{A}_\pi$, both over $\Sigma = \mathcal{A} \times \mathcal{O}$. Let $\mathbb{B} = \mathbb{A} \otimes \mathbb{A}_\pi$ be the stochastic WFA obtained by coupling the environment and the policy and $\rho_{\mathbb{B}} \in \mathcal{P}(\Sigma^\omega)$ the corresponding stochastic process. Suppose that $\mathbb{B}$ is (C,θ)-geometrically mixing. Suppose π satisfies the exploration Assumption 1 with parameter ε. Suppose the importance sampling constant κ in Algorithm 3 satisfies $\kappa\varepsilon > 1$. Let $\tilde{\mathbb{A}}_t = \langle \tilde{\alpha}_t, \beta, \{A_\sigma\}\rangle$ be the WFA defined in Section 5, where the initial vector is $\tilde{\alpha}_t = (1/t)\sum_{s=0}^{t-1}\alpha^\top (A/\kappa)^s$. Let $\bar{\mathbb{A}} = \mathbb{A} \otimes \mathbb{A}_{unif}$ be the stochastic WFA $\langle \alpha, \beta, A_\sigma/|\mathcal{A}|\rangle$ obtained by coupling the environment $\mathbb{A}$ with the uniform random policy. Suppose $\bar{\mathbb{A}}$ is $(\bar{C},\bar{\theta})$-geometrically mixing. Let $L = \max_{w \in \mathcal{U}\cdot\mathcal{V}} |w|$, $\tilde{m} = \sum_{u \in \mathcal{U}, v \in \mathcal{V}} \tilde{f}_t(uv)$, and $\bar{m} = \sum_{u \in \mathcal{U}, v \in \mathcal{V}} \tilde{f}^{unif}_t(uv)$, where $\tilde{f}_t = f_{\tilde{\mathbb{A}}_t}$ and $\tilde{f}^{unif}_t$ is the function computed by the stochastic WFA obtained by Césaro averaging $\bar{\mathbb{A}}$ over t steps. Let $d = \sum_{w \in \mathcal{U}\cdot\mathcal{V}} |w|_{\mathcal{U},\mathcal{V}}$. Then for any $\delta \in (0,1)$ we have*

$$\mathbb{P}\left(\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2 > \sqrt{\frac{\tilde{m}}{t\varepsilon^L(1-\kappa^{-2}\varepsilon^{-2})}} + \sqrt{\frac{2\bar{m}}{t\varepsilon^{2L}}\left(L + \frac{\bar{C}}{1-\bar{\theta}}\right)} + \frac{C}{\theta(1-\theta)\varepsilon^L}\sqrt{\frac{2d\ln(1/\delta)}{t}}\right) \leqslant \delta \ .$$

Proof of Theorem 6:

Let us introduce the function $g(\xi) = \|\widehat{H}^{\mathcal{U},\mathcal{V}}_t - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2$. We first control $\|g\|_{Lip}$ then $\mathbb{E}[g(\xi)]$, before applying Theorem 1.

Step 1: Control of $\|g\|_{Lip}$.

Let $\xi, \xi' \in \Sigma^\omega$ be trajectories $\xi = x_1x_2\cdots$ and $\xi' = x'_1x'_2\cdots$ differing by one element, say at position ℓ. That is, $x_s = x'_s$ for all $s \neq \ell$. We note that

$$\begin{aligned}\left| \|\widehat{H}^{U,V}_{t,\xi} - \tilde{H}^{U,V}_t\|_2 - \|\widehat{H}^{U,V}_{t,\xi'} - \tilde{H}^{U,V}_t\|_2 \right| &\leqslant \|\widehat{H}^{U,V}_{t,\xi} - \widehat{H}^{U,V}_{t,\xi'}\|_2 \\ &\leqslant \sqrt{\sum_{u\in U}\sum_{v\in V}(\widehat{f}_{t,\xi}(uv) - \widehat{f}_{t,\xi'}(uv))^2} \\ &= \frac{1}{t}\sqrt{\sum_{u\in U}\sum_{v\in V}\left(\sum_{s=0}^{t-1}\varphi_{s,uv}(\xi) - \varphi_{s,uv}(\xi')\right)^2} .\end{aligned}$$

Next we take any $w \in \mathcal{U}\cdot\mathcal{V}$ and use $x_i = (o_i, a_i)$ to write

$$\begin{aligned}|\varphi_{s,w}(\xi) - \varphi_{s,w}(\xi')| &= \left| \frac{\mathbb{I}\{o_{s+1}a_{s+1}\cdots o_{s+|uv|}a_{s+|uv|} = w\}}{\kappa^s\pi(a_1\cdots a_{s+|w|}|o_1\cdots o_{s+|w|})} - \frac{\mathbb{I}\{o'_{s+1}a'_{s+1}\cdots o'_{s+|uv|}a'_{s+|uv|} = w\}}{\kappa^s\pi(a'_1\cdots a'_{s+|w|}|o'_1\cdots o'_{s+|w|})} \right| \\ &\leqslant \frac{1}{\kappa^s}\left(\frac{1}{\pi(a_1\cdots a_{s+|w|}|o_1\cdots o_{s+|w|})} + \frac{1}{\pi(a'_1\cdots a'_{s+|w|}|o'_1\cdots o'_{s+|w|})}\right) \\ &\leqslant \frac{2}{\kappa^s\varepsilon^{s+|w|}} ,\end{aligned}$$

where we used the exploration assumption $\pi(u^{\mathcal{A}}|u^{\mathcal{O}}) \geqslant \varepsilon^{|u|}$ for all $u \in \Sigma^\star$.

From the expression above we see that for any $w \in \mathcal{U}\cdot\mathcal{V}$ we have

$$\sum_{s=0}^{t-1}\varphi_{s,w}(\xi) - \varphi_{s,w}(\xi') \leqslant \frac{2}{(1-1/(\kappa\varepsilon))\varepsilon^{|w|}} ,$$

where we used that $\kappa\varepsilon > 1$. Thus, we can conclude that

$$\|g\|_{Lip} \leqslant \frac{2}{t(1-1/(\kappa\varepsilon))}\sqrt{\sum_{w\in\mathcal{U}\cdot\mathcal{V}}\frac{|w|_{\mathcal{U},\mathcal{V}}}{\varepsilon^{2|w|}}} \leqslant \frac{2}{t\varepsilon^L(1-1/(\kappa\varepsilon))}\sqrt{\sum_{w\in\mathcal{U}\cdot\mathcal{V}}|w|_{\mathcal{U},\mathcal{V}}} .$$

Note that $d = \sum_{w\in\mathcal{U}\cdot\mathcal{V}}|w|_{\mathcal{U},\mathcal{V}}$ is the quantity defined in the statement of Theorem 3.

Step 2: Control of $\mathbb{E}[g(\xi)]$. We now want to control the following quantity $\mathbb{E}[\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2]$. We start in the same way as in the proof of Theorem 3.

Step 2.1. By Jensen's inequality, the norm of $\widehat{H}^{\mathcal{U},\mathcal{V}}_t - \tilde{H}^{\mathcal{U},\mathcal{V}}_t$ is controlled by its Frobenius norm

$$\begin{aligned}\mathbb{E}[\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2]^2 &\leqslant \sum_{u\in\mathcal{U}}\sum_{v\in\mathcal{V}}\mathbb{E}\left[\left(\widehat{f}_{t,\xi}(uv) - \tilde{f}_t(uv)\right)^2\right] \\ &= \sum_{w\in\mathcal{U}\cdot\mathcal{V}}|w|_{\mathcal{U},\mathcal{V}}\mathbb{E}\left[\left(\widehat{f}_{t,\xi}(w) - \tilde{f}_t(w)\right)^2\right] .\end{aligned}$$

Recall that in Section 5 we showed that $\mathbb{E}[\widehat{f}_{t,\xi}(w)] = \tilde{f}_t(w)$ for any $w \in \Sigma^\star$. Hence the expression above is a sum of variances, each of which can be written as

$$\mathbb{E}\left[\left(\widehat{f}_{t,\xi}(w) - \tilde{f}_t(w)\right)^2\right] = \mathbb{E}\left[\widehat{f}_{t,\xi}(w)^2\right] - \tilde{f}_t(w)^2 . \tag{8}$$

Now we recall the definitions of the quantities appearing in this expression:

$$\begin{aligned}\widehat{f}_{t,\xi}(w) &= \frac{1}{t}\sum_{s=0}^{t-1}\varphi_{s,w}(\xi) \\ &= \frac{1}{t}\sum_{s=0}^{t-1}\frac{\mathbb{I}\{o_{s+1}a_{s+1}\cdots o_{s+|w|}a_{s+|w|} = w\}}{\kappa^s\pi(a_1\cdots a_{s+|w|}|o_1\cdots o_{s+|w|})} \enspace, \\ \tilde{f}_t(w) &= \frac{1}{t}\sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)}{\kappa^s} \\ &= \frac{1}{t}\sum_{s=0}^{t-1}\alpha^\top\left(\frac{A}{\kappa}\right)^s A_w\beta \enspace .\end{aligned}$$

Therefore, we can expand the squares in (8) as follows:

$$\begin{aligned}\mathbb{E}\left[\widehat{f}_{t,\xi}(w)^2\right] &= \frac{1}{t^2}\left(\sum_{s=0}^{t-1}\mathbb{E}\left[\varphi_{s,w}(\xi)^2\right] + 2\sum_{0\leqslant s<s'\leqslant t-1}\mathbb{E}\left[\varphi_{s,w}(\xi)\varphi_{s',w}(\xi)\right]\right) \enspace, \\ \tilde{f}_t(w)^2 &= \frac{1}{t^2}\left(\sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)^2}{\kappa^{2s}} + 2\sum_{0\leqslant s<s'\leqslant t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)f_{\mathbb{A}}(\Sigma^{s'}w)}{\kappa^{s+s'}}\right) \enspace .\end{aligned}$$

Using these expression we now bound the difference in (8) by considering the "squared" and the "cross" terms separately.

Step 2.2. We start with the "squared" terms and note that for any $0 \leqslant s \leqslant t-1$ and $w \in \mathcal{U}\cdot\mathcal{V}$ we have

$$\begin{aligned}\mathbb{E}\left[\varphi_{s,w}(\xi)^2\right] &= \sum_{w'\in\Sigma^s}\frac{f_{\mathbb{B}}(w'w)}{\kappa^{2s}\pi(w'^{\mathcal{A}}w^{\mathcal{A}}|w'^{\mathcal{O}}w^{\mathcal{O}})^2} \\ &= \sum_{w'\in\Sigma^s}\frac{f_{\mathbb{A}}(w'w)}{\kappa^{2s}\pi(w'^{\mathcal{A}}w^{\mathcal{A}}|w'^{\mathcal{O}}w^{\mathcal{O}})} \\ &\leqslant \frac{f_{\mathbb{A}}(\Sigma^s w)}{\kappa^{2s}\varepsilon^{s+|w|}} \\ &= \frac{f_{\mathbb{A}}(\Sigma^s w)}{\kappa^s(\kappa\varepsilon)^s\varepsilon^{|w|}} \enspace .\end{aligned}$$

Using Cauchy–Schwartz to sum these terms over t we obtain:

$$\begin{aligned}\sum_{s=0}^{t-1}\mathbb{E}\left[\varphi_{s,w}(\xi)^2\right] &\leqslant \sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)}{\kappa^s(\kappa\varepsilon)^s\varepsilon^{|w|}} \\ &\leqslant \frac{1}{(1-1/(\kappa^2\varepsilon^2))\varepsilon^{|w|}}\left(\sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)^2}{\kappa^{2s}}\right)\end{aligned}$$

Using this bound we can now see that the contribution of the "squared" terms to (8) is at most

$$\begin{aligned}\frac{1}{t^2}\left(\sum_{s=0}^{t-1}\mathbb{E}\left[\varphi_{s,w}(\xi)^2\right] - \sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)^2}{\kappa^{2s}}\right) &\leqslant \frac{1}{t^2}\left(\frac{1}{(1-1/(\kappa^2\varepsilon^2))\varepsilon^{|w|}} - 1\right)\left(\sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)^2}{\kappa^{2s}}\right) \\ &\leqslant \frac{1}{t^2(1-1/(\kappa^2\varepsilon^2))\varepsilon^{|w|}}\left(\sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)^2}{\kappa^{2s}}\right) \enspace .\end{aligned}$$

This expression can be further simplified by noting that $\varepsilon \leqslant 1/|\mathcal{A}|$ implies $\kappa > |\mathcal{A}|$ and therefore $f_{\mathbb{A}}(\Sigma^s w)/\kappa^s \leqslant f_{\mathbb{A}}(\Sigma^s w)/|\mathcal{A}|^s \leqslant 1$ since this corresponds to the probability of observing $w^{\mathcal{O}}$ when taking the actions in $w^{\mathcal{A}}$ after the first s actions have been chosen by a uniform random policy. Thus, we get

$$\frac{1}{t^2}\left(\sum_{s=0}^{t-1}\mathbb{E}\left[\varphi_{s,w}(\xi)^2\right] - \sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)^2}{\kappa^{2s}}\right) \leqslant \frac{1}{t^2(1-1/(\kappa^2\varepsilon^2))\varepsilon^{|w|}}\left(\sum_{s=0}^{t-1}\frac{f_{\mathbb{A}}(\Sigma^s w)}{\kappa^s}\right)$$
$$= \frac{\tilde{f}_t(w)}{t(1-1/(\kappa^2\varepsilon^2))\varepsilon^{|w|}} \enspace .$$

To complete this step we sum this bound for all $w \in \mathcal{U}\cdot\mathcal{V}$ to control the contribution of the "squared" terms in (8):

$$\sum_{w\in\mathcal{U}\cdot\mathcal{V}} |w|_{\mathcal{U},\mathcal{V}}\frac{\tilde{f}_t(w)}{t(1-1/(\kappa^2\varepsilon^2))\varepsilon^{|w|}} \leqslant \frac{1}{t(1-1/(\kappa^2\varepsilon^2))\varepsilon^{L}}\sum_{w\in\mathcal{U}\cdot\mathcal{V}}|w|_{\mathcal{U},\mathcal{V}}\tilde{f}_t(w)$$
$$= \frac{1}{t(1-1/(\kappa^2\varepsilon^2))\varepsilon^{L}}\sum_{u\in\mathcal{U},v\in\mathcal{V}}\tilde{f}_t(uv) \enspace ,$$

where $L = \max_{w\in\mathcal{U}\cdot\mathcal{V}}|w|$.

Step 2.3. We now focus on controlling the "cross" terms in (8) of the form

$$\mathbb{E}\left[\varphi_{s,w}(\xi)\varphi_{s',w}(\xi)\right] - \frac{f_{\mathbb{A}}(\Sigma^s w)f_{\mathbb{A}}(\Sigma^{s'} w)}{\kappa^{s+s'}} \enspace . \tag{9}$$

Using the same notation $\Sigma_w^{s'-s} = w\Sigma^{s'-s}\cap\Sigma^{s'-s}w$ as in the proof of Theorem 3, we first note that

$$\mathbb{E}\left[\varphi_{s,w}(\xi)\varphi_{s',w}(\xi)\right] = \sum_{x\in\Sigma^s\Sigma_w^{s'-s}}\frac{f_{\mathbb{B}}(x)}{\kappa^{s+s'}\pi(x^{\mathcal{A}}_{1:s+|w|}|x^{\mathcal{O}}_{1:s+|w|})\pi(x^{\mathcal{A}}_{1:s'+|w|}|x^{\mathcal{O}}_{1:s'+|w|})}$$
$$= \sum_{x\in\Sigma^s\Sigma_w^{s'-s}}\frac{f_{\mathbb{A}}(x)}{\kappa^{s+s'}\pi(x^{\mathcal{A}}_{1:s+|w|}|x^{\mathcal{O}}_{1:s+|w|})}$$
$$\leqslant \sum_{x\in\Sigma^s\Sigma_w^{s'-s}}\frac{f_{\mathbb{A}}(x)}{\kappa^{s+s'}\varepsilon^{s+|w|}}$$
$$= \frac{f_{\mathbb{A}}(\Sigma^s\Sigma_w^{s'-s})}{\kappa^{s+s'}\varepsilon^{s+|w|}}$$
$$= \frac{\alpha^\top A^s A_w^{s'-s}\beta}{\kappa^{s+s'}\varepsilon^{s+|w|}} \enspace ,$$

where we used the notation $A_w^{s'-s} = \sum_{x\in\Sigma_w^{s'-s}} A_x$. We also define $\tilde{A} = A/\kappa$ and $\alpha_s^\top = \alpha^\top\tilde{A}^s$. Then we can write (9) as

$$\frac{\alpha^\top A^s A_w^{s'-s}\beta}{\kappa^{s+s'}\varepsilon^{s+|w|}} - \frac{(\alpha^\top A^s A_w\beta)(\alpha^\top A^{s'}A_w\beta)}{\kappa^{s+s'}} = \alpha_s^\top\left(\frac{A_w^{s'-s}}{\kappa^{s'}\varepsilon^{s+|w|}} - A_w\beta\alpha_{s'}^\top A_w\right)\beta \enspace . \tag{10}$$

To bound this quantity we proceed by considering two cases.

Step 2.4. First suppose that $s'-s \geqslant l = |w|$. In this case we have $A_w^{s'-s} = A_w A^{s'-s-l}A_w$ and (10) equals to

$$\alpha_s^\top A_w\left(\frac{A^{s'-s-l}}{\kappa^{s'}\varepsilon^{s+l}} - \beta\alpha_{s'}^\top\right)A_w\beta = \alpha_s^\top A_w\left(\frac{\tilde{A}^{s'-s-l}}{\kappa^{s+l}\varepsilon^{s+l}} - \beta\alpha_{s+l}^\top\tilde{A}^{s'-s-l}\right)A_w\beta \enspace .$$

Now we apply the same argument we used to bound the "cross" terms in the case of stochastic WFA using cone norms. In particular, we consider the stochastic WFA $\bar{\mathbb{A}} = \langle \alpha, \beta, \bar{A}_\sigma \rangle$, where $\bar{A}_\sigma = A_\sigma / |\mathcal{A}|$. Note this is the stochastic WFA obtained by coupling environment $\mathbb{A}$ with the random policy that at each step chooses each action independently with probability $1/|\mathcal{A}|$. Now we let $\| \cdot \|_\beta$ and $\| \cdot \|_{\beta,\star}$ denote the cone norms corresponding to $\bar{\mathbb{A}}$. Using Lemma 3 we see that the following hold for all $w \in \Sigma^\star$:

$$\|A_w \beta\|_\beta = |\mathcal{A}|^l \|\bar{A}_w \beta\|_\beta \leqslant |\mathcal{A}|^l$$
$$\|\alpha_s^\top A_w\|_{\beta,\star} = \frac{|\mathcal{A}|^{s+l}}{\kappa^s} \|\alpha^\top \bar{A}^s \bar{A}_w\|_{\beta,\star} = \frac{|\mathcal{A}|^{s+l}}{\kappa^s} \alpha^\top \bar{A}^s \bar{A}_w \beta \ ,$$

where we used the notation $\bar{A} = A/|\mathcal{A}|$. We also note that for any vector satisfying $\|u\|_{\beta,\star} \leqslant 1$ we have

$$\|u^\top \beta \alpha_s^\top\|_{\beta,\star} \leqslant \|\alpha_s^\top\|_{\beta,\star} = \frac{|\mathcal{A}|^s}{\kappa^s} \|\alpha^\top \bar{A}^s\|_{\beta,\star} \leqslant \frac{|\mathcal{A}|^s}{\kappa^s} \leqslant \frac{1}{\kappa^s \varepsilon^s} \ .$$

This last bound can now be combined with the argument used in the case of stochastic WFA to show that

$$\begin{aligned}
\left\| \frac{\tilde{A}^{s'-s-l}}{\kappa^{s+l}\varepsilon^{s+l}} - \beta \alpha_{s+l}^\top \tilde{A}^{s'-s-l} \right\|_\beta &= \sup_{\|u\|_{\beta,\star} \leqslant 1} \left\| u^\top \frac{\tilde{A}^{s'-s-l}}{\kappa^{s+l}\varepsilon^{s+l}} - u^\top \beta \alpha_{s+l}^\top \tilde{A}^{s'-s-l} \right\|_{\beta,\star} \\
&\leqslant \sup_{\|u_1\|_{\beta,\star} \leqslant 1} \sup_{\|u_2\|_{\beta,\star} \leqslant 1} \left\| u_1^\top \frac{\tilde{A}^{s'-s-l}}{\kappa^{s+l}\varepsilon^{s+l}} - u_2^\top \frac{\tilde{A}^{s'-s-l}}{\kappa^{s+l}\varepsilon^{s+l}} \right\|_{\beta,\star} \\
&= \frac{|\mathcal{A}|^{s'-s+l}}{\kappa^{s'}\varepsilon^{s+l}} \sup_{\|u_1\|_{\beta,\star} \leqslant 1} \sup_{\|u_2\|_{\beta,\star} \leqslant 1} \left\| u_1^\top \bar{A}^{s'-s-l} - u_2^\top \bar{A}^{s'-s-l} \right\|_{\beta,\star} \\
&\leqslant \frac{|\mathcal{A}|^{s'-s+l}}{\kappa^{s'}\varepsilon^{s+l}} \mu^{\bar{\mathbb{A}}}_{s'-s-l} \ ,
\end{aligned}$$

where we used the definition of the mixing coefficient $\mu^{\bar{\mathbb{A}}}_{s'-s-l}$ for stochastic WFA $\bar{\mathbb{A}}$.

We now observe that $|\mathcal{A}| \leqslant 1/\varepsilon < \kappa$ implies $|\mathcal{A}|^{s'+l}/\kappa^{s+s'}\varepsilon^{s+l} \leqslant 1/\kappa^s \varepsilon^{s+2l}$. Finally, by plugging all these bounds together on an application of Hölder's inequality yields:

$$\left| \alpha_s^\top A_w \left(\frac{A^{s'-s-l}}{\kappa^{s'}\varepsilon^{s+l}} - \beta \alpha_{s'}^\top \right) A_w \beta \right| \leqslant \frac{\mu^{\bar{\mathbb{A}}}_{s'-s-l}}{\kappa^s \varepsilon^{s+2l}} \alpha^\top \bar{A}^s \bar{A}_w \beta \ .$$

Step 2.5. Now we consider the case $s' - s < l = |w|$. Using the fact that this implies $\Sigma_w^{s'-s} \subset w\Sigma^{s'-s}$, then

$$\alpha_s^\top A_w^{s'-s} \beta \leqslant \alpha_s^\top A_w A^{s'-s} \beta = |\mathcal{A}|^{s'-s} \alpha_s^\top A_w \bar{A}^{s'-s} \beta = |\mathcal{A}|^{s'-s} \alpha_s^\top A_w \beta \ ,$$

where we used $\bar{A}\beta = \beta$. Therefore, we can bound the expression in (10) as

$$\begin{aligned}
\alpha_s^\top \left(\frac{A_w^{s'-s}}{\kappa^{s'}\varepsilon^{s+l}} - A_w \beta \alpha_{s'}^\top A_w \right) \beta \leqslant \alpha_s^\top A_w \beta \left(\frac{|\mathcal{A}|^{s'-s}}{\kappa^{s'}\varepsilon^{s+l}} - \alpha_{s'}^\top A_w \beta \right) &\leqslant \frac{|\mathcal{A}|^{s'-s}}{\kappa^{s'}\varepsilon^{s+l}} \alpha_s^\top A_w \beta \\
&= \frac{|\mathcal{A}|^{s'+l}}{\kappa^{s'+s}\varepsilon^{s+l}} \alpha^\top \bar{A}^s \bar{A}_w \beta \leqslant \frac{1}{\kappa^s \varepsilon^{s+2l}} \alpha^\top \bar{A}^s \bar{A}_w \beta \ .
\end{aligned}$$

Step 2.6. Finally, we can combine the bounds above by summing over all $w \in \mathcal{U} \cdot \mathcal{V}$ and all $0 \leqslant s < s' \leqslant t-1$ in the same way we did for PFA. We first note that from Steps 2.4 and 2.5 we obtain the following bound for (10):

$$\alpha_s^\top \left(\frac{A_w^{s'-s}}{\kappa^{s'} \varepsilon^{s+|w|}} - A_w \beta \alpha_{s'}^\top A_w \right) \beta \leqslant \frac{\bar{f}_s(w)}{\kappa^s \varepsilon^{s+2|w|}} \left(\mu^{\bar{\mathbb{A}}}_{s'-s-|w|} \mathbb{I}\{s'-s \geqslant |w|\} + \mathbb{I}\{s'-s < |w|\} \right) \ .$$

Now let $l = |w|$ and note that $\mu^{\bar{\mathbb{A}}}_{s'-s-l} \leqslant \bar{C}\bar{\theta}^{s'-s-l}$, where $\bar{C}$ and $\bar{\theta}$ are the geometric mixing constants for stochastic WFA $\bar{\mathbb{A}}$. Thus, summing first over s' we get

$$\sum_{s'=s+1}^{t-1} \mu^{\bar{\mathbb{A}}}_{s'-s-|w|} \mathbb{I}\{s'-s \geqslant |w|\} + \mathbb{I}\{s'-s < |w|\} \leqslant l + \frac{\bar{C}}{1-\bar{\theta}} \ .$$

Therefore, writing $\mathcal{W}_l$ for all words of length l in $\mathcal{W} = \mathcal{U} \cdot \mathcal{V}$ we get:

$$\begin{aligned}
&\frac{2}{t^2} \sum_{w \in \mathcal{U}\cdot\mathcal{V}} |w|_{\mathcal{U},\mathcal{V}} \sum_{0 \leqslant s < s' \leqslant t-1} \left(\mathbb{E}\left[\varphi_{s,w}(\xi)\varphi_{s',w}(\xi)\right] - \frac{f_{\mathbb{A}}(\Sigma^s w) f_{\mathbb{A}}(\Sigma^{s'} w)}{\kappa^{s+s'}} \right) \\
&\leqslant \frac{2}{t^2} \sum_{l=0}^{\infty} \sum_{w \in \mathcal{W}_l} \frac{|w|_{\mathcal{U},\mathcal{V}}}{\varepsilon^{2l}} \left(l + \frac{\bar{C}}{1-\bar{\theta}} \right) \sum_{s=0}^{t-2} \frac{\bar{f}_s(w)}{\kappa^s \varepsilon^s} \\
&\leqslant \frac{2}{t} \sum_{l=0}^{\infty} \frac{1}{\varepsilon^{2l}} \left(l + \frac{\bar{C}}{1-\bar{\theta}} \right) \sum_{w \in \mathcal{W}_l} |w|_{\mathcal{U},\mathcal{V}} \sum_{s=0}^{t-1} \frac{\bar{f}_s(w)}{t} \\
&\leqslant \frac{2}{t\varepsilon^{2L}} \left(L + \frac{\bar{C}}{1-\bar{\theta}} \right) \sum_{u \in \mathcal{U}, v \in \mathcal{V}} \tilde{f}_t^{unif}(uv) \ ,
\end{aligned}$$

where we used that $\kappa\varepsilon > 1$ and $\tilde{f}_t^{unif}(w) = (1/t)\sum_{s=0}^{t-1} \bar{f}_s(w)$.

Step 2.7. Our final bound for $\mathbb{E}[\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2]$ is now obtained by combining the results from Step 2.2 and 2.6:

$$\mathbb{E}[\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2]^2 \leqslant \frac{1}{t\varepsilon^L(1-1/(\kappa\varepsilon)^2)} \sum_{u \in \mathcal{U}, v \in \mathcal{V}} \tilde{f}_t(uv) + \frac{2}{t\varepsilon^{2L}} \left(L + \frac{\bar{C}}{1-\bar{\theta}} \right) \sum_{u \in \mathcal{U}, v \in \mathcal{V}} \tilde{f}_t^{unif}(uv) \ .$$

Note that $\tilde{m} = \sum_{u \in \mathcal{U}, v \in \mathcal{V}} \tilde{f}_t(uv)$ and $\bar{m} = \sum_{u \in \mathcal{U}, v \in \mathcal{V}} \tilde{f}_t^{unif}(uv)$ are the quantities defined in the statement of Theorem 6.

Step 3. Application of Theorem 1 It follows directly from Theorem 1 that with probability at least $1-\delta$ we have

$$\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2 \leqslant \mathbb{E}[\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2] + \eta_{\rho_{\mathbb{B}}} \|g\|_{Lip} \sqrt{\frac{t\ln(1/\delta)}{2}} \ .$$

Using that $\rho_{\mathbb{B}}$ is (C,θ)-geometrically mixing and Lemma 4 we can bound the η-mixing coefficient as $\eta_{\rho_{\mathbb{B}}} \leqslant C/(\theta(1-\theta))$. Thus, by plugging our estimates for $\|g\|_{Lip}$ and $\mathbb{E}[\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2]$ we obtain that with probability at least $1-\delta$:

$$\|\widehat{H}^{\mathcal{U},\mathcal{V}}_{t,\xi} - \tilde{H}^{\mathcal{U},\mathcal{V}}_t\|_2 \leqslant \sqrt{\frac{\tilde{m}}{t\varepsilon^L(1-\kappa^{-2}\varepsilon^{-2})}} + \sqrt{\frac{2\bar{m}}{t\varepsilon^{2L}} \left(L + \frac{\bar{C}}{1-\bar{\theta}} \right)} + \frac{C}{\theta(1-\theta)\varepsilon^L} \sqrt{\frac{2d\ln(1/\delta)}{t}} \ .$$

□
