

440 **A Proof**

441 **A.1 Proof of Performance Difference Distinction via State Sequences**

442 Following the previous work [19], our analysis will make use of the discounted future state distribution,  
 443  $d^\pi$ , which is defined as

$$d^\pi(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P(s_t = s | \pi, \mathcal{M})$$

444 It allows us to express the expected discounted total reward compactly as

$$\begin{aligned} J(\pi) &= \sum_{t=0}^{\infty} \gamma^t E_{s_t, a_t, s_{t+1}} [R(s_t, a_t, s_{t+1}) | \pi, \mathcal{M}] \\ &= \sum_{t=0}^{\infty} \gamma^t \int_{\mathcal{S}} R_\pi(s) P(s_t = s | \pi, \mathcal{M}) ds \\ &= \int_{\mathcal{S}} R_\pi(s) \sum_{t=0}^{\infty} \gamma^t P(s_t = s | \pi, \mathcal{M}) ds \tag{1} \\ &= \frac{1}{1 - \gamma} \int_{\mathcal{S}} R_\pi(s) d_\pi(s) ds \\ &= \frac{1}{1 - \gamma} E_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [R(s, a, s')], \tag{2} \end{aligned}$$

445 where we define  $R_\pi(s) := E_{a \sim \pi, s' \sim P} [R(s, a, s')]$ . It should be clear from  $a \sim \pi(\cdot | s)$  and  $s' \sim$   
 446  $P(\cdot | s, a)$  that  $a$  and  $s'$  depend on  $s$ . Thus, the reward function  $R_\pi$  is only related to  $s$  when the policy  
 447  $\pi$  is fixed.

448 Firstly, we prove that the distance between two state sequence distributions obtained from two distinct  
 449 policies serves as an upper bound on the performance difference between those policies, provided  
 450 that certain assumptions regarding the reward function hold.

451 **Theorem 1.** *Suppose that the reward function  $R(s, a, s') = R(s)$  is related to the state  $s$ , then the*  
 452 *performance difference between two arbitrary policies  $\pi_1$  and  $\pi_2$  is bounded by the L1 norm of the*  
 453 *difference between their state sequence distributions:*

$$|J(\pi_1) - J(\pi_2)| \leq \frac{R_{max}}{1 - \gamma} \cdot \|P(s_0, s_1, s_2, \dots | \pi_1, \mathcal{M}) - P(s_0, s_1, s_2, \dots | \pi_2, \mathcal{M})\|_1, \tag{3}$$

454 where  $P(s_0, s_1, s_2, \dots | \pi_1, \mathcal{M})$  means the joint distribution of the infinite-horizon state sequence  
 455  $\mathbf{S} = \{s_0, s_1, s_2, \dots\}$  conditioned on the policy  $\pi$  and the environment model  $\mathcal{M}$ .

456 *Proof.* According to the equation (1), the difference in performance between two policies  $\pi_1, \pi_2$  can  
 457 be bounded as follows.

$$\begin{aligned} |J(\pi_1) - J(\pi_2)| &\leq R_{max} \cdot \sum_{t=0}^{\infty} \gamma^t \int_{\mathcal{S}} |P(\mathbf{s}_t = s | \pi_1, \mathcal{M}) - P(\mathbf{s}_t = s | \pi_2, \mathcal{M})| ds \\ &\leq R_{max} \cdot \sum_{t=0}^T \gamma^t \int_{\mathcal{S}} \left| \int_{\mathcal{S}^T} P(s_0, \dots, s_{t-1}, s, s_{t+1}, \dots, s_T | \pi_1, \mathcal{M}) \right. \\ &\quad \left. - P(s_0, \dots, s_{t-1}, s, s_{t+1}, \dots, s_T | \pi_2, \mathcal{M}) ds_0 \cdots ds_{t-1} ds_{t+1} \cdots ds_T \right| ds \\ &\quad + R_{max} \cdot 2 \sum_{t=T+1}^{\infty} \gamma^t, \quad \forall T \geq 1 \end{aligned}$$

$$\begin{aligned}
&\leq R_{\max} \sum_{t=0}^T \gamma^t \int_{\mathcal{S}^{T+1}} |P(s_0, \dots, s_T | \pi_1, \mathcal{M}) - P(s_0, \dots, s_T | \pi_2, \mathcal{M})| ds_0 \cdots ds_T \\
&\quad + R_{\max} \cdot 2 \sum_{t=T+1}^{\infty} \gamma^t, \quad \forall T \geq 1 \\
&= \frac{R_{\max}}{1-\gamma} \cdot \int_{\mathcal{S}^{T+1}} |P(s_0, \dots, s_T | \pi_1, \mathcal{M}) - P(s_0, \dots, s_T | \pi_2, \mathcal{M})| ds_0 \cdots ds_T \\
&\quad + R_{\max} \cdot 2 \sum_{t=T+1}^{\infty} \gamma^t, \quad \forall T \geq 1.
\end{aligned}$$

458 Let  $T \rightarrow \infty$ , then we obtain the bound proposed by (3).  $\square$

459 We are further interested in bounding the performance difference between two policies by their  
460 state sequences in the frequency domain. Benefiting from the properties of the discrete-time Fourier  
461 transform (DTFT), we can constrain the performance difference using the Fourier transform over the  
462 interval  $[0, 2\pi]$ , instead of using the distribution functions of the state sequences in unbounded space.

463 **Theorem 2.** *Suppose that  $\mathcal{S} \subset \mathbb{R}^D$  the reward function  $R(s, a, s') = R(s)$  is an  $n$ th-degree  
464 polynomial function with respect to  $s \in \mathcal{S}$ , then for any two policies  $\pi_1$  and  $\pi_2$ , their performance  
465 difference can be bounded as follows:*

$$|J(\pi_1) - J(\pi_2)| \leq \frac{\sqrt{D}}{1-\gamma} \cdot \sum_{k=1}^n \frac{\|R^{(k)}(0)\|_D}{k!} \cdot \max_{1 \leq i \leq D} \sup_{\omega_i \in [0, 2\pi]} |F_{\pi_1}^{(k)}(\omega_i) - F_{\pi_2}^{(k)}(\omega_i)|, \quad (4)$$

466 where  $F_{\pi}^{(k)}(\omega)$  denotes the DTFT of the time series  $\mathbf{S}^{(k)} = \{\mathbf{s}_0^k, \mathbf{s}_1^k, \mathbf{s}_2^k, \dots\}$  for any integer  
467  $k \in [1, n]$  and  $\mathbf{S}^{(k)}$  means the  $k$ th power of the state sequence produced by the policy  $\pi$ . The  
468 dimensionality of  $\omega$  is the same as  $s$ .

469 *Proof.* For sake of simplicity, we define  $p_t(s|\pi_i) = P(\mathbf{s}_t = s|\pi_i, \mathcal{M})$  for  $i = 1, 2$ . We denote  $\varepsilon_t$  as

$$\varepsilon_t = \int_{\mathcal{S}} R(s) [p_t(s|\pi_1) - p_t(s|\pi_2)] ds. \quad (5)$$

470 Based on the Taylor series expansion, we can rewrite the reward function as  $R(s) = \sum_{k=0}^n \frac{R^{(k)}(0)^{\top}}{k!} s^k$ ,  
471 then for any integer  $k \in [1, n]$ , we have

$$\begin{aligned}
|\varepsilon_t| &\leq \sum_{k=0}^n \frac{\|R^{(k)}(0)\|_D}{k!} \cdot \left\| \int_{\mathcal{S}} [s^k p_t(s|\pi_1) - s^k p_t(s|\pi_2)] ds \right\|_D \\
&= \sum_{k=0}^n \frac{\|R^{(k)}(0)\|_D}{k!} \left\| \mathbb{E}_{s \sim p_t(\cdot|\pi_1)} [s^k] - \mathbb{E}_{s \sim p_t(\cdot|\pi_2)} [s^k] \right\|_D. \quad (6)
\end{aligned}$$

472 Since the inverse DTFT of  $F_{\pi}^{(k)}(\omega)$  is the original time series  $\mathbf{S}^{(k)}$ , we have

$$\mathbb{E}_{s_i \sim p_t(\cdot|\pi)} [s_i^k] = \frac{1}{2\pi} \int_0^{2\pi} F_{\pi}^{(k)}(\omega_i) e^{j\omega_i t} d\omega_i, \quad \forall i = 1, 2, \dots, D. \quad (7)$$

473 Then we have

$$\begin{aligned}
\left| \mathbb{E}_{s_i \sim p_t(\cdot|\pi_1)} [s_i^k] - \mathbb{E}_{s_i \sim p_t(\cdot|\pi_2)} [s_i^k] \right| &\leq \frac{1}{2\pi} \int_0^{2\pi} |F_{\pi_1}^{(k)}(\omega_i) - F_{\pi_2}^{(k)}(\omega_i)| \cdot |e^{j\omega_i t}| d\omega_i \\
&\leq \sup_{\omega_i \in [0, 2\pi]} |F_{\pi_1}^{(k)}(\omega_i) - F_{\pi_2}^{(k)}(\omega_i)|. \quad (8)
\end{aligned}$$

474 Substituting (8) into (6), then we obtain

$$|\varepsilon_t| \leq \sqrt{D} \cdot \sum_{k=1}^n \frac{\|R^{(k)}(0)\|_D}{k!} \cdot \max_{1 \leq i \leq D} \sup_{\omega_i \in [0, 2\pi]} |F_{\pi_1}^{(k)}(\omega_i) - F_{\pi_2}^{(k)}(\omega_i)|.$$



508 Note that for  $k < d$ ,  $x_k$  is the eigenvector corresponding to  $\lambda_k$ . Since the column vectors of  $X$  are  
 509 linearly dependent, there exist  $\vec{c} = [c_0, c_1, \dots, c_{D-1}]$  not all zero, such that  $\mu_0 = \sum_{k=0}^{D-1} c_k \vec{x}_k = X\vec{c}$ .  
 510 Thus, we have

$$P^n \mu_0 = \sum_{k=0}^{d-1} c_k \lambda_k^n \vec{x}_k + \sum_{k=d}^{D-1} c_k P^n \vec{x}_k. \quad (9)$$

511 For any Jordan cell  $J_k$ , let  $\alpha_k$  be the multiplicity of  $\lambda_k$ , then

$$J_k^n = \begin{bmatrix} \lambda_k & 1 & & & \\ & \lambda_k & 1 & & \\ & & \ddots & \ddots & \\ & & & \lambda_k & 1 \\ & & & & \lambda_k \end{bmatrix}_{\alpha_k \times \alpha_k}^n = \begin{bmatrix} \lambda_k^n & C_n^{n-1} \lambda_k^{n-1} & \dots & C_n^{n-\alpha_k+1} \lambda_k^{n-\alpha_k+1} \\ & \lambda_k^n & \dots & C_n^{n-\alpha_k+2} \lambda_k^{n-\alpha_k+2} \\ & & \ddots & \vdots \\ & & & \lambda_k^n \end{bmatrix}.$$

512 Since  $\alpha_k$  is fixed for matrix  $P$ , we have  $\lim_{n \rightarrow \infty} J_k^n = \mathbf{0}$  for each  $k = d, \dots, D-1$ . Then the limiting  
 513 vector of (9), denoted by  $P^\infty \mu_0$ , satisfies:

$$P^\infty \mu_0 = \lim_{n \rightarrow \infty} X J^n X^{-1} X \vec{c} = \lim_{n \rightarrow \infty} \sum_{k=0}^{d-1} c_k \lambda_k^n \vec{x}_k = \lim_{n \rightarrow \infty} \mu^{(n)},$$

514 where we denote  $\mu^{(n)} = \sum_{k=0}^{d-1} c_k (e^{j \frac{2\pi k}{d}})^n \vec{x}_k$ . Let  $r = n \pmod{d}$ , then we have

$$\mu^{(n)} = \mu^{(r)} = \sum_{k=0}^{d-1} c_k (\xi^k)^r \vec{x}_k, \quad \forall n \geq 1.$$

515 Therefore, the probability sequence  $\{P^n \mu_0\}_{n \geq 1}$  will split into  $d$  converging subsequences and has  $d$   
 516 cyclic limiting probability distributions when  $n \rightarrow \infty$ , denoted as

$$\mu_\infty^r = \sum_{k=0}^{d-1} c_k (\xi^k)^r \vec{x}_k, \quad r = 0, 1, \dots, d-1.$$

517 Thus,  $P^n \mu_0$  is asymptotically periodic with period  $d$  if  $d > 1$  and asymptotically aperiodic if  
 518  $d = 1$ .  $\square$

519 We now consider a more general state space that may not necessarily be irreducible. According to  
 520 the Decomposition theorem of the Markov chain [28], the finite state space  $S$  can be partitioned  
 521 uniquely as a set of transient states and one or several irreducible closed sets of recurrent states.  
 522 According to [29], after performing an appropriate permutation of rows and columns, we can rewrite  
 523 the transition probability matrix  $P$  in its canonical form:

$$P = \left[ \begin{array}{cccc|c} R_1 & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & R_2 & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & R_\alpha & \mathbf{0} \\ \hline T_1 & T_2 & \dots & T_\alpha & Q \end{array} \right],$$

524 where  $R_1, \dots, R_\alpha$  represent the probability submatrices corresponding to the recurrent classes,  $Q$   
 525 represents the probability submatrix corresponding to the transient states, and  $T_1, \dots, T_\alpha$  represent  
 526 the probability submatrices corresponding to the transitions between transient and recurrent classes  
 527  $R_1, \dots, R_\alpha$  respectively.

528 **Theorem 3.** Suppose that the state space  $S$  is finite with a transition probability matrix  $P \in \mathbb{R}^{|S| \times |S|}$   
 529 and  $S$  has  $\alpha$  recurrent classes. Let  $R_1, R_2, \dots, R_\alpha$  be the probability submatrices corresponding  
 530 to the recurrent classes and let  $d_1, d_2, \dots, d_\alpha$  be the number of the eigenvalues of modulus 1 that  
 531 the submatrices  $R_1, R_2, \dots, R_\alpha$  has. Then for any initial distribution  $\mu_0$ ,  $P^n \mu_0$  is asymptotically  
 532 periodic with period  $d = \text{lcm}(d_1, d_2, \dots, d_\alpha)$  when  $d > 1$  and asymptotically aperiodic when  $d = 1$ .

533 *Proof.* Since  $P$  is a block upper-triangular, it can be shown that the eigenvalues of  $P$  are equal to  
 534 the union of the eigenvalues of the diagonal blocks  $R_1, \dots, R_\alpha, Q$ . Note that the  $n$ th-power of  $P$   
 535 satisfies the following expression:

$$P^n = \left[ \begin{array}{cccc|c} R_1^n & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & R_2^n & \cdots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & R_\alpha^n & \mathbf{0} \\ \hline T_1^{(n)} & T_2^{(n)} & \cdots & T_\alpha^{(n)} & Q^n \end{array} \right],$$

536 where  $T_r^{(n)}$  is related to the  $(n-1)$ -th or the lower power of  $R_r$  and  $Q$ . From Theorem 4.3 of [29],  
 537 we obtain that  $\lim_{n \rightarrow \infty} Q^n = \mathbf{0}$ , which implies that all eigenvalues of  $Q$  have modulus less than 1.

538 On the other hand, note that the sum of every row in matrix  $R_r$  is equal to 1, which means  $\lambda = 1$  is  
 539 an eigenvalue of  $R_r$  and all eigenvalues of  $R_r$  satisfy  $|\lambda| \leq 1$ . Thus, the spectral radius of  $P$  is equal  
 540 to 1.

541 Note that the proof of Lemma 1 implies that the asymptotic periodicity of  $P^n \mu_0$  depends on the  
 542 eigenvalues of  $P$  that have modulus 1. Since  $R_r$  is non-negative irreducible with spectral radius 1,  
 543 based on the Perron-Frobenius theorem used in Lemma 1, we can express the eigenvalues of  $R_r$  in  
 544 modulus 1 as:

$$\lambda_{r,k} = e^{j \frac{2\pi k}{d_r}}, \quad k = 0, 1, \dots, d_r - 1.$$

545 Based on the above discussion, it is easy to check that  $\bigcup_{r=1}^{\alpha} \{\lambda_{r,0}, \dots, \lambda_{r,d_r-1}\}$  is the set of all  
 546 eigenvalues of modulus 1 of  $P$ . Rewrite  $P$  in its Jordan canonical form  $P = X J X^{-1}$ , where

$$J = \left[ \begin{array}{cccccccc} \lambda_{1,0} & & & & & & & \\ & \ddots & & & & & & \\ & & \lambda_{1,d_1-1} & & & & & \\ & & & \lambda_{2,0} & & & & \\ & & & & \ddots & & & \\ & & & & & \lambda_{\alpha,d_\alpha-1} & & \\ & & & & & & J_{d_1+\dots+d_\alpha} & \\ & & & & & & & \ddots \\ & & & & & & & & J_s \end{array} \right]$$

547 and  $X = [\vec{x}_0, \vec{x}_1, \dots, \vec{x}_{D-1}]$  is an invertible matrix. Similar to the proof in Lemma 1, we get

$$P^\infty \mu_0 = \lim_{n \rightarrow \infty} \sum_{r=1}^{\alpha} \sum_{k=0}^{d_r-1} c_k (e^{j \frac{2\pi k}{d_r}})^n \vec{x}_k := \lim_{n \rightarrow \infty} \mu^{(n)}.$$

548 Let  $d = \text{lcm}(d_1, d_2, \dots, d_\alpha)$  and  $r = n \pmod{d}$ , then we have

$$\mu^{(n)} = \mu^{(r)}, \quad \forall n \geq 1.$$

549 Therefore, the probability sequence  $\{P^n \mu_0\}_{n \geq 1}$  will split into  $d$  converging subsequences and has  $d$   
 550 cyclic limiting probability distributions when  $n \rightarrow \infty$ , denoted as

$$\mu_\infty^r = \sum_{r=1}^{\alpha} \sum_{k=0}^{d_r-1} c_k e^{j \frac{2\pi k r}{d_r}} \vec{x}_k, \quad r = 0, 1, \dots, d-1.$$

551 Thus,  $P^n \mu_0$  is asymptotically periodic with period  $d$  if  $d > 1$  and asymptotically aperiodic if  $d = 1$ .  
 552 This completes the proof.  $\square$

### 553 A.3 Proof of the Convergence of Our Auxiliary Loss

554 In this section, we provide a detailed derivation of the learning objective of SPF. As the DTFT of  
 555 discrete-time state sequences is a continuous function that is difficult to compute, we practically  
 556 sample the DTFT at  $L$  equally-spaced points.

$$[\mathcal{F}\tilde{s}_t]_k = \sum_{n=0}^{+\infty} [\tilde{s}_t]_n e^{-j\frac{2\pi k}{L}n}, \quad k = 0, 1, \dots, L-1. \quad (10)$$

557 As a result, the prediction target takes the form of a matrix with dimensions of  $L * D$ , where  $D$  denotes  
 558 the dimension of the state space. The auxiliary task is designed to encourage the representation  
 559 to predict the Fourier transform of the state sequences using the current state-action pair as input.  
 560 Specifically, we define the prediction target  $F_{\pi,p}(s_t, a_t)$  as follows:

$$F_{\pi,p}(s_t, a_t) = \mathcal{F}\tilde{s}(s_t, a_t) = \left\{ \sum_{n=0}^{+\infty} [\tilde{s}(s_t, a_t)]_n e^{-j\frac{2\pi k}{L}n} \right\}_{k=0}^{L-1}, \quad (11)$$

561 For simplicity of notation, we substitute  $F(s_t, a_t)$  for  $F_{\pi,p}(s_t, a_t)$  in the following. We can derive  
 562 that the DTFT functions at successive time steps are related to each other in a recursive form:

$$\begin{aligned} [F(s_t, a_t)]_k &= \sum_{n=0}^{+\infty} \gamma^n \cdot e^{-j\frac{2\pi k}{L}n} \cdot E_{\pi,p} [s_{t+n+1} | s_t = s, a_t = a] \\ &= E_p [s_{t+1} | s_t = s, a_t = a] + \gamma \cdot e^{-j\frac{2\pi k}{L}} \cdot \\ &\quad E_{s_{t+1} \sim p, a_{t+1} \sim \pi} \left[ \sum_{n=0}^{+\infty} \gamma^n \cdot e^{-j\frac{2\pi k}{L}n} \cdot E_p [s_{t+n+2} | s_{t+1}, a_{t+1}] \right] \\ &= [\tilde{s}_t]_0 + \gamma \cdot e^{-j\frac{2\pi k}{L}} \cdot E_{\pi,p} [[F(s_{t+1}, a_{t+1})]_k], \quad \forall k = 0, 1, \dots, L-1. \end{aligned}$$

563 We can further express the above equation as a matrix-form recursive formula as follows:

$$F(s_t, a_t) = \tilde{\mathbf{S}}_t + \Gamma E_{\pi,p} [F(s_{t+1}, a_{t+1})], \quad (12)$$

564 where

$$\tilde{\mathbf{S}}_t = [[\tilde{s}_t]_0, \dots, [\tilde{s}_t]_0]^T \in \mathbb{R}^{L \times D},$$

565

$$\Gamma = \gamma \begin{bmatrix} 1 & & & & \\ & e^{-j\frac{2\pi}{L}} & & & \\ & & e^{-j\frac{4\pi}{L}} & & \\ & & & \ddots & \\ & & & & e^{-j\frac{(L-1)\pi}{L}} \end{bmatrix}.$$

566 Similar to the TD-learning of value functions, we can prove that the above recursive relationship (12)  
 567 can be reformulated as a contraction mapping  $\mathcal{T}$ . Due to the properties of contraction mappings, we  
 568 can iteratively apply the operator  $\mathcal{T}$  to compute the target DTFT function until convergence in tabular  
 569 settings.

570 **Theorem 4.** Let  $\mathcal{F}$  denote the set of all functions  $F : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{C}^{L \times D}$  and define the norm on  $\mathcal{F}$  as

$$\|F\|_{\mathcal{F}} := \sup_{\substack{s \in \mathcal{S} \\ a \in \mathcal{A}}} \max_{0 \leq k < L} \|[F(s, a)]_k\|_D,$$

571 where  $[F(s, a)]_k$  represents the  $k$ th row vector of  $F(s, a)$ . We show that the mapping  $\mathcal{T} : \mathcal{F} \rightarrow \mathcal{F}$   
 572 defined as

$$\mathcal{T}F(s_t, a_t) = \tilde{\mathbf{S}}_t + \Gamma E_{\pi,p} [F(s_{t+1}, a_{t+1})] \quad (13)$$

573 is a contraction mapping, where  $\tilde{\mathbf{S}}_t$  and  $\Gamma$  are defined as above.

574 *Proof.* For any  $F_1, F_2 \in \mathcal{F}$ , we have

$$\begin{aligned}
\|\mathcal{T}F_1 - \mathcal{T}F_2\|_{\mathcal{F}} &= \sup_{\substack{s \in \mathcal{S} \\ a \in \mathcal{A}}} \max_{0 \leq k < L} \left\| s + \gamma e^{-j \frac{2\pi k}{K}} E_{\substack{s' \sim P(\cdot|s,a) \\ a' \sim \pi(\cdot|s')}} \left[ [F_1(s', a')]_k | s, a \right] \right. \\
&\quad \left. - s - \gamma e^{-j \frac{2\pi k}{K}} E_{\substack{s' \sim P(\cdot|s,a) \\ a' \sim \pi(\cdot|s')}} \left[ [F_2(s', a')]_k | s, a \right] \right\|_D \\
&\leq \gamma \cdot \max_{0 \leq k < L} \sup_{\substack{s \in \mathcal{S} \\ a \in \mathcal{A}}} \left\| E_{\substack{s' \sim P(\cdot|s,a) \\ a' \sim \pi(\cdot|s')}} \left[ [F_1(s', a')]_k - [F_2(s', a')]_k | s, a \right] \right\|_D \\
&\leq \gamma \cdot \max_{0 \leq k < L} \sup_{\substack{s' \in \mathcal{S} \\ a' \in \mathcal{A}}} \|[F_1(s', a') - F_2(s', a')]_k\|_D \\
&= \gamma \cdot \|F_1 - F_2\|_{\mathcal{F}}.
\end{aligned}$$

575 Note that  $\gamma \in [0, 1)$ , which implies that  $\mathcal{T}$  is a contraction mapping.  $\square$

## 576 B Pseudo-code of SPF

The training procedure of SPF is shown in the pseudo-code as follows:

---

### Algorithm 1 State Sequences Prediction via Fourier Transform (SPF)

---

Denote parameters of the online encoder  $(\phi_s, \phi_{s,a})$ , predictor  $\mathcal{F}$ , and projection  $\psi$  as  $\theta_{\text{aux}}$   
Denote parameters of the target encoder  $(\hat{\phi}_s, \hat{\phi}_{s,a})$ , predictor  $\hat{\mathcal{F}}$ , and projection  $\hat{\psi}$  as  $\hat{\theta}_{\text{aux}}$   
Denote parameters of actor model  $\pi$  and critic model  $Q$  for RL agents as  $\theta_{\text{RL}}$   
Denote the smoothing coefficient and update interval for target network updates as  $\tau$  and  $K$   
Initialize replay buffer  $\mathcal{D}$  and parameters  $\theta_{\text{aux}}, \theta_{\text{RL}}$   
**for** each environment step  $t$  **do**  
     $a_t \sim \pi(\cdot | \phi_s(s_t))$   
     $s_{t+1}, r_{t+1} \sim p(\cdot | s_t, a_t)$   
     $\mathcal{D} \leftarrow \mathcal{D} \cup (s_t, a_t, s_{t+1}, r_{t+1})$   
    sample a minibatch of  $\{(s_t, a_t, s_{t+1}, r_{t+1})\}$  from  $\mathcal{D}$   
     $\theta_{\text{aux}} \leftarrow \theta_{\text{aux}} - \alpha_{\text{aux}} \nabla_{\theta_{\text{aux}}} L_{\text{pred}}(\theta_{\text{aux}}, \hat{\theta}_{\text{aux}})$   
    resampling a minibatch of  $\{(s_t, a_t, s_{t+1}, r_{t+1})\}$  from  $\mathcal{D}$   
     $\bar{s}_t \leftarrow \phi_s(s_t)$   
     $z_{s_t, a_t} \leftarrow \phi_{s,a}(\phi_s(s_t), a_t)$   
    update the RL agent parameters  $\theta_{\text{RL}}$  with the representations  $\bar{s}_t, z_{s_t, a_t}$   
    update parameters of target networks with  $\hat{\theta}_{\text{aux}} \leftarrow \tau \theta_{\text{aux}} + (1 - \tau) \hat{\theta}_{\text{aux}}$  every  $K$  steps  
**end for**

---

577

## 578 C Network Details

579 The encoders  $\phi_s$  and  $\phi_{s,a}$  share the same architecture. Each layer of the encoders uses MLP-  
580 DenseNet [16], a slightly modified version of DenseNet. For each MuJoCo task, the incremental  
581 number of hidden units per layer is selected from  $\{30, 40\}$ , while the number of layers is selected  
582 from  $\{6, 8\}$  (see Table 1). Both the predictor  $\mathcal{F}$  and the projection  $\psi$  apply a 2-layer MLP. We  
583 divide the last layer of the predictor into two heads as the real part  $\mathcal{F}_{\text{Re}}$  and the imaginary part  $\mathcal{F}_{\text{Im}}$ ,  
584 respectively, since the prediction target of our auxiliary task is complex-valued. With respect to the  
585 projection module, we add an additional 2-layer MLP (referred to as *Projection2*) after the original  
586 online projection to perform a dimension-invariant nonlinear transformation on the predicted DTFT  
587 that has been projected to a lower-dimensional space. We do not apply this nonlinear operation to the  
588 target projection. This additional step is carried out to prevent the projection from collapsing to a  
589 constant value in the case where the online and target projections share the same architecture.

590 In Fourier analysis, the low-frequency components of the DTFT contain information about the  
591 long-term trends of the signal, with higher signal energy, while the high-frequency components of the

Table 1: Detailed setting of the encoder for six MuJoCo tasks.

| Environment    | Number of Layers | Number of Units per Layer | Activation Function |
|----------------|------------------|---------------------------|---------------------|
| HalfCheetah-v2 | 8                | 30                        | Swish               |
| Walker2d-v2    | 6                | 40                        | Swish               |
| Hopper-v2      | 6                | 40                        | Swish               |
| Ant-v2         | 6                | 40                        | Swish               |
| Swimmer-v2     | 6                | 40                        | Swish               |
| Humanoid-v2    | 8                | 40                        | Swish               |

592 DTFT reflect the amount of short-term variation present in the state sequences. Therefore, we attempt  
593 to preserve the overall information of the low and high-frequency components of the predicted DTFT  
594 by directly computing the cosine similarity distance without undergoing the dimensionality reduction  
595 process. For the remaining frequency components of the predicted DTFT, we first utilize projection  
596 layers to perform dimensionality reduction, followed by calculating the cosine similarity distance.  
597 The sum of these three distances is used as the final loss function, which we call *freqloss*.

## 598 D Hyperparameters

Table 2: Hyperparameters of auxiliary prediction tasks.

| Hyperparameter                                      | Setting            |
|---|--------------------|
| Optimizer   | Adam               |
| Discount $\gamma$                                   | 0.99               |
| Learning rate                                       | 0.0003             |
| Number of batch size                                | 256                |
| Predictor: Number of hidden layers                  | 1                  |
| Predictor: Number of hidden units per layer         | 1024               |
| Predictor: Activation function                      | ReLU               |
| Projection: Number of hidden layers                 | 1                  |
| Projection: Number of hidden units per layer        | 512                |
| Projection: Activation function                     | ReLU               |
| Projection2: Number of hidden layers                | 1                  |
| Projection2: Number of hidden units per layer       | 512                |
| Projection2: Activation function                    | ReLU               |
| Number of discrete points for sampling the DTFT $L$ | 128                |
| The dimensionality of the output of projection      | 512                |
| Replay buffer size                                  | 100,000            |
| Pre-training steps                                  | 10000              |
| Target smoothing coefficient $\tau$                 | 0.01               |
| Target update interval $K$                          | 1000               |
| <i>Hyperparameters of SPF-SAC</i>                   |                    |
| Each module: Normalization Layer                    | BatchNormalization |
| Random collection steps before pre-training         | 10,000             |
| <i>Hyperparameters of SPF-PPO</i>                   |                    |
| Each module: Normalization Layer                    | LayerNormalization |
| Random collection steps before pre-training         | 4,000              |
| $\theta_{\text{aux}}$ update interval $K_2$         |                    |
| HalfCheetah-v2                                      | 5                  |
| Walker2d-v2   | 2                  |
| Hopper-v2   | 150                |
| Ant-v2  | 150                |
| Swimmer-v2  | 200                |
| Humanoid-v2   | 1                  |

599 We select  $L = 128$  as the number of discrete points sampled over one period of DTFT. In practice, due  
600 to the symmetry conjugate of DTFT, the predictor  $\mathcal{F}$  only predicts  $\frac{L}{2} + 1$  points on the left half of our  
601 frequency map, as mentioned in Section 5.2. The projection module described in Section 5.3 projects  
602 the predicted value, a matrix with the dimension of  $L * D$ , into a 512-dimensional vector. To update  
603 target networks, we overwrite the target network parameters with an exponential moving average of  
604 the online network parameters, with a smoothing coefficient of  $\tau = 0.01$  for every  $K = 1000$  steps.

605 In order to eliminate dependency on the initial parameters of the policy, we use a random policy to  
606 collect transitions into the replay buffer [30] for the first 10K time steps for SAC, and 4K time steps  
607 for PPO. We also pretrain the representations with the aforementioned random collected samples to  
608 stabilize inputs to each RL algorithm, as described in [16].

609 The network architectures, optimizers, and hyperparameters of SAC and PPO are the same as those  
610 used in their original papers, except that we use mini-batches of size 256 instead of 100. As for PPO,  
611 we perform  $K_2$  gradient updates of  $\theta_{\text{aux}}$  for every  $K_2$  steps of data sampling. The update interval  $K_2$   
612 is set differently for six MuJoCo tasks and can be found in Table 2.

## 613 E Visualization

614 To demonstrate that the representations learned by SPF effectively capture the structural information  
615 contained in infinite-step state sequences, we compare the true state sequences with the states  
616 recovered from the predicted DTFT via the inverse DTFT.

617 Specifically, we first generate a state sequence from the trained policy and select a goal state  $s_t$  at a  
618 certain time step. Next, we choose a historical state  $s_{t-k}$  located  $k$  steps past the goal state and select  
619 an action  $a_{t-k}$  based on the trained policy  $\pi(\cdot|s_{t-k})$  as the inputs of our trained predictor. We then  
620 obtain the DTFT  $F_{t-k} := F_{\pi}(s_{t-k}, a_{t-k})$  of state sequences starting from the state  $s_{t-k+1}$ . Next,  
621 we compute the  $k$ th element of the inverse DTFT of  $F_{t-k}$  and obtain a recovered state  $\hat{s}_t$ , which  
622 represents that we predict the future goal state using the historical state located  $k$  steps past the goal  
623 state. By selecting a sequence of states over a specific time interval as the goal states and repeating  
624 the aforementioned procedures, we will obtain a state sequence recovered by  $k$ -step prediction. In  
625 Figure 5(b), 6(b), 7(b), 8(b), 9(b) and 10(b), we visualize the true state sequence (the blue line) and  
626 the recovered state sequences (the red lines) via  $k$ -step predictions for  $k = 1, 2, 3, 4, 5$ . Note that  
627 the lighter red line corresponds to predictions made by historical states from a more distant time  
628 step. We conduct the visualization experiment on six MuJoCo tasks using the representations and  
629 predictors trained by SPF-SAC or SPF-PPO. Due to the large dimensionality of the states in Ant-v2  
630 and Humanoid-v2, which contain many zero values, we have chosen to visualize only six dimensions  
631 of their states, respectively. The fine distinctions between the true state sequences and the recovered  
632 state sequences from our trained representations and predicted FT indicates that our representation  
633 effectively captures the inherent structures of future state sequences.

634 Furthermore, we provide a visualization that compares the true DTFT and the predicted DTFT in  
635 Figure 5(a), 6(a), 7(a), 8(a), 9(a) and 10(a). To accomplish this, we use our trained policies to interact  
636 with the environments and select the state sequences of the 200 last steps of an episode. The blue  
637 lines represent the true DTFT of these state sequences, while the orange line represents the predicted  
638 DTFT using the online encoder and predictor trained by our learned policies. It is evident that the true  
639 DTFT and the predicted DTFT exhibit significant differences. These results demonstrate the ability  
640 of SPF to effectively extract the underlying structural information in infinite-step state sequences  
641 without relying on high prediction accuracy.

## 642 F Code

643 Codes for the proposed method are available at [https://anonymous.4open.science/r/spf\\_nips\\_2023-10D1/README.md](https://anonymous.4open.science/r/spf_nips_2023-10D1/README.md).  
644

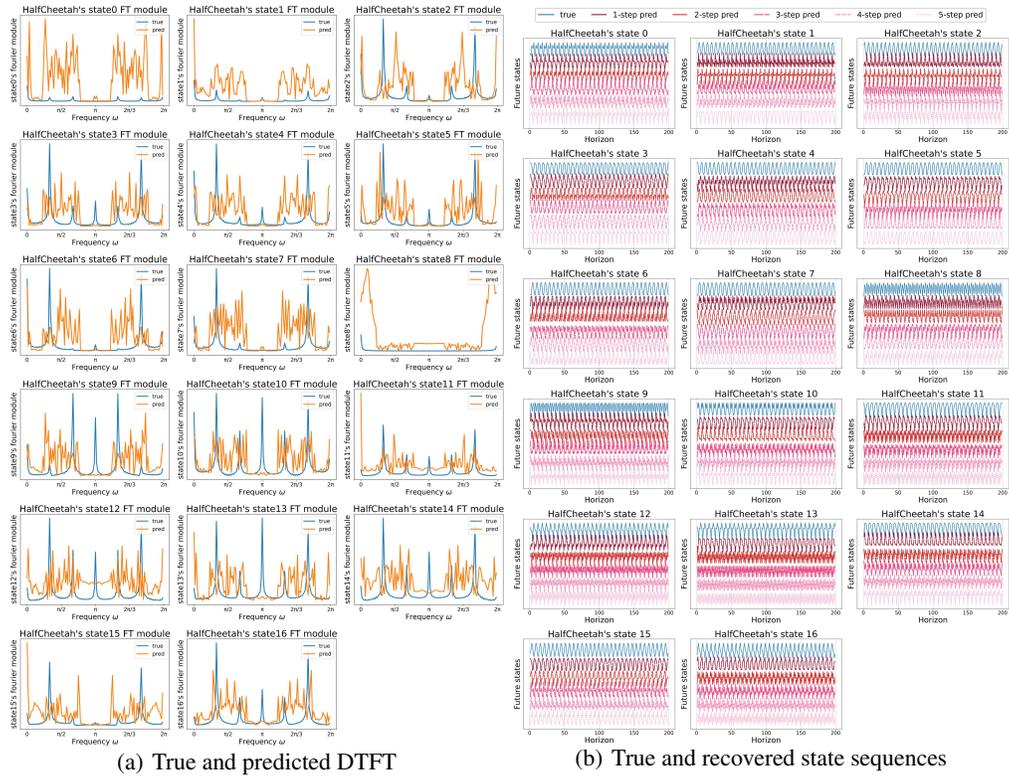


Figure 5: Predicted values via representations trained by SPF-SAC on HalfCheetah-v2

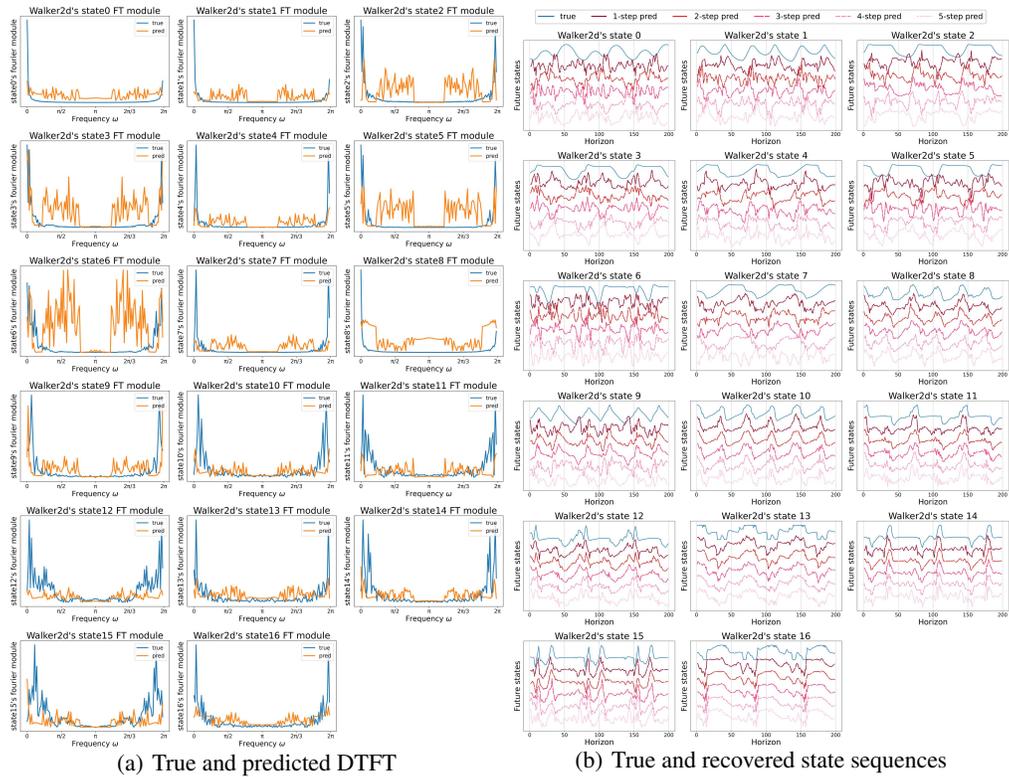


Figure 6: Predicted values via representations trained by SPF-SAC on Walker2d-v2

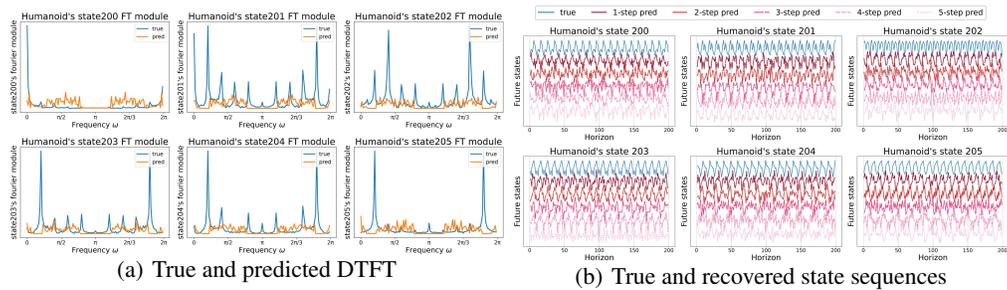


Figure 7: Predicted values via representations trained by SPF-SAC on Humanoid-v2

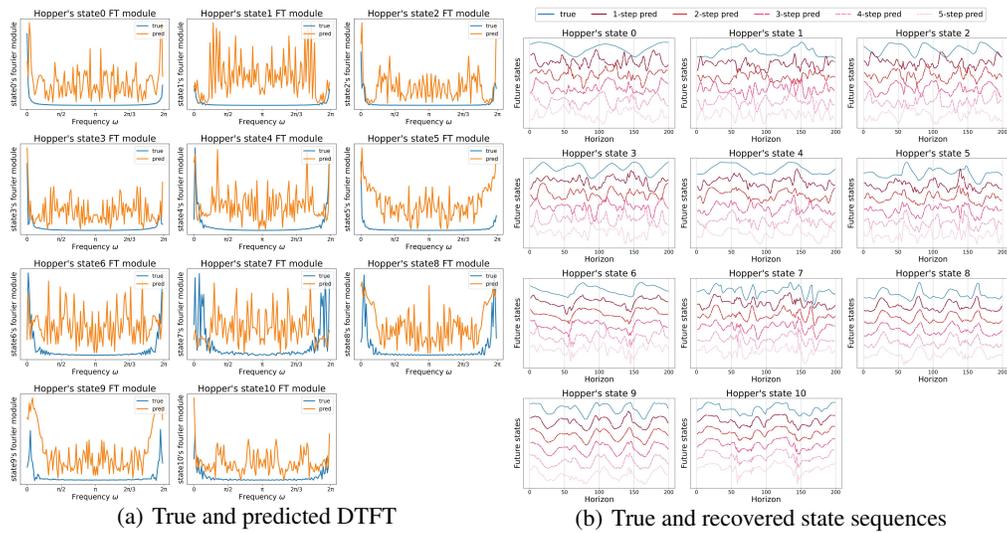


Figure 8: Predicted values via representations trained by SPF-PPO on Hopper-v2

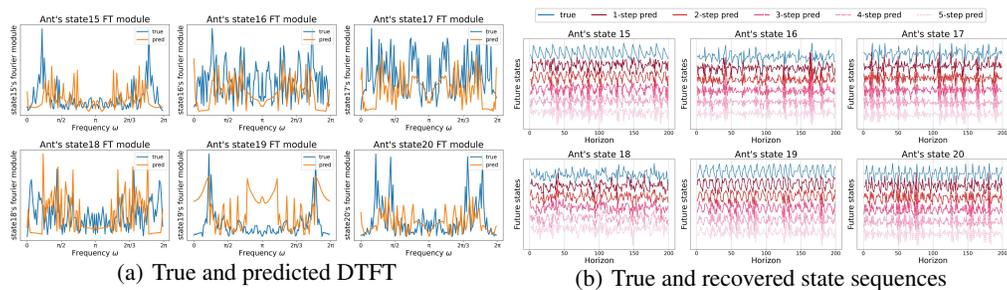


Figure 9: Predicted values via representations trained by SPF-PPO on Ant-v2

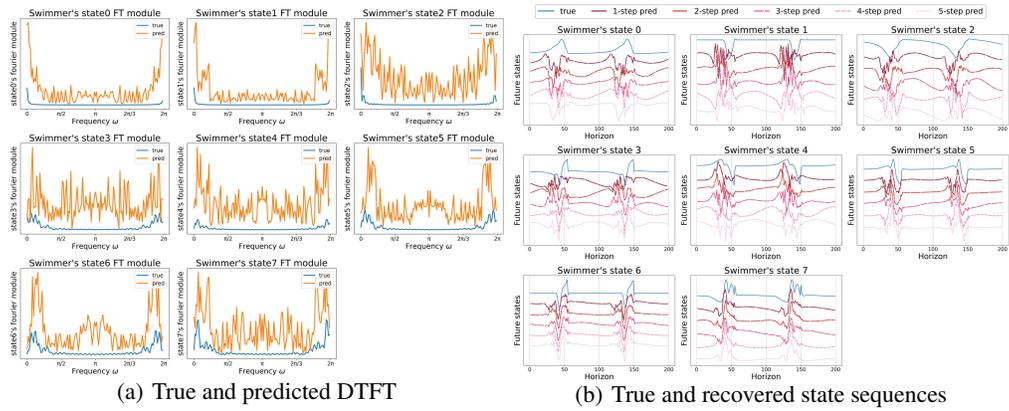


Figure 10: Predicted values via representations trained by SPF-PPO on Swimmer-v2