Proceedings of Machine Learning Research vol 132:1–23, 2021 32nd International Conference on Algorithmic Learning Theory

Learning and Testing Irreducible Markov Chains via the k-cover Time

Siu On Chan SIUON@CSE.CUHK.EDU.HK
Qinghua Ding QHDING@CSE.CUHK.EDU.HK
Sing Hei Li SHLI@CSE.CUHK.EDU.HK
The Chinese University of Hong Kong

Editors: Vitaly Feldman, Katrina Ligett and Sivan Sabato

Abstract

We give a unified way of testing and learning finite Markov chains from a single Markovian trajectory, using the idea of k-cover time introduced here. The k-cover time is the expected length of a random walk to cover every state at least k times. This generalizes the notion of cover time in the literature. The error metric in the testing and learning problems is the infinity matrix norm between the transition matrices, as considered by Wolfer and Kontorovich.

Specifically, we show that if we can learn or test discrete distributions using k samples, then we can learn or test Markov chains using a number of samples equal to the k-cover time of the chain, up to constant factors. We then derive asymptotic bounds on the k-cover time in terms of the number of states, minimum stationary probability and the cover time of the chain. Our bounds are tight for reversible Markov chains and almost tight (up to logarithmic factors) for irreducible ones.

Our results on k-cover time yield sample complexity bounds for a wider range of learning and testing tasks (including learning, uniformity testing, identity testing, closeness testing and their tolerant versions) over Markov chains, and can be applied to a broader family of Markov chains (irreducible and reversible ones) than previous results which only applies to ergodic ones.

Keywords: Markov chains, learning and testing

1. Introduction

Learning and testing discrete distributions is an active research area (see, e.g., Anthony and Bartlett (2009); Batu et al. (2001); Chan et al. (2014) and the references therein). Classical results include $\Theta(n/\epsilon^2)$ as sample complexity for learning Anthony and Bartlett (2009) and $\Theta(\sqrt{n}/\epsilon^2)$ as sample complexity for uniformity testing Paninski (2008). A number of other learning and testing problems have been proposed and studied as well, including identity testing, closeness testing and tolerant learning/testing (see, e.g., the survey by Canonne Canonne (2017)).

We consider these problems when the samples are not iid, but instead generated from a finite Markov chain, as considered in Daskalakis et al. (2017); Wolfer and Kontorovich (2019b, 2020a). Following Wolfer and Kontorovich (2019b, 2020a), we use the infinity matrix norm as the distance measure.[1] The main challenge in the Markovian case is that, since the samples are dependent, the mixing properties of the chain needs to be taken into consideration.

Consider a Markov chain over discrete state space $[n] = \{1, 2, ..., n\}$. Given the initial state X_0, one can generate the Markovian trajectory $X_1, X_2, ..., X_T$ according to the transition probabilities $\mathbb{P}(X_t = j|X_{t-1} = i) = p_{ij}$ for all $t \geq 1$. Denote by $M = (p_{ij})_{i,j\in[n]}$ the *transition matrix* of this chain. The Markov chain is *irreducible* if for all $i, j \in [n]$, there exists some $t \in \mathbb{N}$ such that $(M^t)_{ij} > 0$. For each irreducible Markov chain, the fundamental theorem of Markov chain

1. Note that the infinity matrix norm is equivalent to the metric $|||\cdot|||$ in Wolfer and Kontorovich (2019b).

guarantees a unique stationary distribution $\pi = (\pi_1, ..., \pi_n) \in \Delta^{n-1}$, which is entry-wise positive, such that $\pi M = \pi$. Here $\Delta^{n-1} \triangleq \{\pi : \mathbf{1}^T \pi = 1, \pi \in \mathbb{R}_+^n\}$ is the $(n-1)$-dimensional probability simplex. We denote the *minimum stationary probability* as $\pi_* \triangleq \min_{i \in [n]} \pi_i$.

If the chain is *reversible* in addition to being irreducible, it then satisfies the *detailed balance* condition: $p_{ij}\pi_i = p_{ji}\pi_j, \forall i, j \in [n]$. If a Markov chain is reversible, then the eigenvalues of its transition matrix M are all real and can be denoted as $\lambda_1 = 1 > \lambda_2 \geq ... \geq \lambda_n \geq -1$. The *spectral gap* of this chain is $\gamma \triangleq 1 - \lambda_2$, and the *absolute spectral gap* of this chain is $\gamma_* \triangleq \min\{1 - |\lambda_2|, 1 - |\lambda_n|\}$. It is well known that γ_* characterizes the mixing time t_{mix} of reversible chains via the inequalities $\Omega(1/\gamma_*) \leq t_{\text{mix}} \leq O(\ln(1/\pi_*)/\gamma_*)$.

(Uniformly) ergodic chains form a sub-family of irreducible chains that also satisfies the aperiodicity condition. For ergodic chains, the mixing time is similarly characterized by Paulin's *pseudo-spectral gap* γ_{ps} Paulin (2015). This quantity generalizes the absolute spectral gap by suitably *reversiblizing* the chain. Formally, $\gamma_{\text{ps}} \triangleq \max_{k \geq 1} \frac{1}{k} \gamma((M^T)^k M^k)$.

Given a Markovian trajectory $X^m = (X_0, \ldots, X_m)$ from some unknown Markov chain M up to time m, we are interested in learning M from this trajectory. A popular choice in the literature is the plug-in estimator $\hat{M}$ defined as $\hat{M} = (N_{ij}/m)_{i,j \in [n]}$, where N_{ij} is the the number of transitions from state i to state j in this trajectory. The quality of any estimator $\hat{M}$ is then valued by its closeness to M under some distance measure $d(M, \hat{M})$.

Besides learning, there are also testing tasks including uniformity testing, identity testing and closeness testing. We list the following four natural learning and testing tasks for Markov chains here.

1. **(ϵ, δ)-Learning :** Given small constants $\delta, \epsilon \in (0, 1)$, and a Markovian trajectory X_1^m from some unknown chain M, an (ϵ, δ)-learning algorithm $\mathcal{A}$ outputs a transition matrix $\hat{M} = \mathcal{A}(X_1^m, n)$ such that $d(\hat{M}, M) \leq \epsilon$ with probability $\geq 1 - \delta$.

2. **(ϵ, δ)-Uniformity Testing :** Given small constants $\delta, \epsilon \in (0, 1)$, and a Markovian trajectory X_1^m from some unknown chain M, an (ϵ, δ)-uniformity testing algorithm $\mathcal{A}(X_1^m, M, n)$ outputs "Yes" if $M = M_u$ and "No" if $d(M, M_u) \geq \epsilon$ with probability $\geq 1 - \delta$. Here $M_u = \frac{1}{n}\mathbf{1}^T\mathbf{1}$ yields exactly uniform i.i.d samples.

3. **(ϵ, δ)-Identity Testing :** Given small constants $\delta, \epsilon \in (0, 1)$, a known reference Markov chain M and a Markovian trajectory X_1^m from another unknown chain M', an (ϵ, δ)-identity testing algorithm $\mathcal{A}(X_1^m, M, n)$ outputs "Yes" if $M = M'$ and "No" if $d(M, M') \geq \epsilon$ with probability $\geq 1 - \delta$.

4. **(ϵ, δ)-Closeness Testing :** Given small constants $\delta, \epsilon \in (0, 1)$, two Markovian trajectories X_1^m, Y_1^m from unknown Markov chains M, M' respectively, an (ϵ, δ)-closeness testing algorithm $\mathcal{A}(X_1^m, Y_1^m, n)$ outputs "Yes" if $M = M'$ and "No" if $d(M, M') \geq \epsilon$ with probability $\geq 1 - \delta$.

For testing problems, there are also tolerant versions: Given $0 < \epsilon_1 < \epsilon_2 < 1$, decide whether $d(M, M') \leq \epsilon_1$ or $d(M, M') \geq \epsilon_2$. These tolerant testing tasks are in general harder than vanilla testing tasks. Details about tolerant testing will be covered in Section 3.

Previous works considered various distance measures $d(M, \hat{M})$: matrix norms, Hellinger-based distance and the minimax prediction risk. We now discuss these distance measures.

- **Infinity Matrix Norm**: Learning Markov chains under the infinity matrix norm $\|\hat{M} - M\|_\infty$ is studied in Wolfer and Kontorovich (2019b). It is shown that a certain estimator (not the empirical one) achieves near-optimal sample complexity $\tilde{\Theta}\left(1/\gamma_{\text{ps}}\pi_* + n/\pi_*\epsilon^2\right)$ for learning ergodic chains. And later, they considered identity testing ergodic chains under this distance, showing that one can achieve near optimal sample complexity $\tilde{\Theta}\left(1/\gamma_{\text{ps}}\pi_* + \sqrt{n}/\pi_*\epsilon^2\right)$ Wolfer and Kontorovich (2020a). Recently, this distance is also studied in Wolfer and Kontorovich (2020b) for learning a Markov chain with a countable state space.

- **A Hellinger-based Distance**: The distance $d_{\surd}(\hat{M}, M)$ was proposed to study identity testing problem of Markov chains in Daskalakis et al. (2017); Cherapanamjeri and Bartlett (2019). However, identity testing under this distance only works for symmetric Markov chains, which is a quite restricted sub-family of Markov chains. Also, this distance measure fails to satisfy the triangle inequality and is not a metric Daskalakis et al. (2017). Thus we do not study learning and testing problems under this distance.

- **Minimax Prediction Risk**: The problem of learning Markov chains under some smooth f-divergence based minimax prediction risk $\rho(\hat{M}, M)$ was studied in Hao et al. (2018). They deduced the near-optimality of the (smoothed) plug-in estimator for achieving low risk, so long as $\min_{i,j} p_{ij} > 0$. This is a fairly strong restriction on Markov chains. Hence, we are not interested in this measure either.

As the above discussion shows, both the Hellinger-based distance and the one based on the minimax population risk put stringent conditions on the families of Markov chains we can study. Thus, we stick with using matrix norms as the distance measure. Moreover, we find the infinity matrix norm $\|\cdot\|_\infty$ natural for its intimate connection to learning and testing with i.i.d. samples, as shown previously in Wolfer and Kontorovich (2019b, 2020a,b). Formally, we have

$$\|M - M'\|_\infty = \max_{i\in[n]} \left\{ \sum_{j\in[n]} |p_{ij} - p'_{ij}| \right\} = \max_{i\in[n]} 2d_{\text{TV}}(\mathbf{p}_i, \mathbf{p}'_i).$$

Here $\mathbf{p}_i = (p_{i1}, p_{i2}, ..., p_{in})$ denotes the outgoing transition probabilities from state i, and d_{TV} denotes the total variation distance.

In this paper, we shed light on the connection between learning and testing problems for Markov chains and those with i.i.d. samples. Specifically, we prove that the sample complexity of learning and testing Markov chains is controlled by a combinatorial quantity $t_{\text{cov}}^{(k)}$ of the unknown chain which we dub as the k-cover time. Informally speaking, we show that if the sample complexity of (ϵ, δ)-learning/testing discrete distributions under d_{TV} is $k(\epsilon, \delta)$, then the sample complexity of learning and testing Markov chains under $\|\cdot\|_\infty$ is upper bounded by $t_{\text{cov}}^{k(\epsilon,\delta')}$ of the unknown chain.

Our argument works for a large family of learning and testing tasks including learning, uniformity testing, identity testing, closeness testing and related tolerant versions of testing problems. Our main results (Theorem 20 and Theorem 27) generalize previous Markov chain learning Wolfer and Kontorovich (2019b) and Markov chain Identity Testing Wolfer and Kontorovich (2020a) results to *every* similarly-defined learning and testing problems on Markov chains. Further, previous results by Wolfer and Kontorovich only hold for ergodic chains, while our results hold more generally for irreducible chains — arguably the most general family of chains having a finite sample complexity guarantee.

Towards our main results, we also prove tight bounds for the k-cover time in terms of k, minimum stationary probability and the cover time. For reversible chains, our bounds $t_{\text{cov}}^{(k)} = \Theta(k/\pi_* + t_{\text{cov}})$ are tight up to constant factors (Theorem 13 and Theorem 19). For irreducible chains, our upper bound $t_{\text{cov}}^{(k)} = \tilde{O}(k/\pi_* + t_{\text{cov}})$ is tight up to logarithmic factors (Theorem 13 and Theorem 22).

2. Preliminaries

In this section, we review some related definitions, lemmas and theorems which will be useful in our analysis. Specifically, we review some backgrounds on testing and learning discrete distributions as well as the Ray–Knight's isomorphism theorem.

2.1. Testing and Learning Discrete Distributions

Testing and learning discrete distributions with i.i.d. samples is a well studied topic, especially under the total variation distance. The following theorem summarizes some results in this area, including sample complexity bounds for learning, identity testing, closeness testing and so on.

Theorem 1 ((ϵ, δ)-learning/testing discrete distributions) *The sample complexity of (ϵ, δ)-learning/testing discrete distributions via i.i.d. samples over state space $[n]$ is as the following.*

1. *(ϵ, δ)-learning (Anthony and Bartlett (2009)): The sample complexity is $\Theta_\delta(n/\epsilon^2)$.*
2. *(ϵ, δ)-uniform testing (Paninski (2008)): The sample complexity is $\Theta_\delta(\sqrt{n}/\epsilon^2)$.*
3. *(ϵ, δ)-identity testing (Batu et al. (2001); Valiant and Valiant (2017)): The sample complexity is $\Theta_\delta(\sqrt{n}/\epsilon^2)$.*
4. *(ϵ, δ)-closeness testing (Chan et al. (2014), Theorem 1): The sample complexity is $\Theta_\delta(\max\{\sqrt{n}/\epsilon^2, n^{2/3}/\epsilon^{4/3}\})$.*
5. *$(\epsilon_1, \epsilon_2, \delta)$-tolerant-uniformity testing (Valiant and Valiant (2011), Theorem 3 and 4): The sample complexity is $O_\delta(n/\ln n(\epsilon_2 - \epsilon_1)^2)$.*
6. *$(\epsilon_1, \epsilon_2, \delta)$-tolerant-identity testing (Valiant and Valiant (2011), Theorem 3 and 4): The sample complexity is $O_\delta(n/\ln n(\epsilon_2 - \epsilon_1)^2)$.*
7. *$(\epsilon_1, \epsilon_2, \delta)$-tolerant-closeness testing (Valiant and Valiant (2011), Theorem 3 and 4): The sample complexity is $O_\delta(n/\ln n(\epsilon_2 - \epsilon_1)^2)$.*
8. *$(\epsilon/2\sqrt{n}, \epsilon, \delta)$-tolerant-uniform testing (Goldreich and Ron (2011), rephrased): The sample complexity is $O_\delta(\sqrt{n}/\epsilon^4)$.*
9. *$(\epsilon^3/300\sqrt{n}\ln n, \epsilon, \delta)$-tolerant-identity testing (Batu et al. (2001), Theorem 24): The sample complexity is $O_\delta(\sqrt{n}\ln n/\epsilon^6)$.*

Here Θ_δ hides the logarithmic term in δ.

In the next section, we will show how Markov chain problems are related to i.i.d. sample problems via the k-cover time.

2.2. Ray–Knight's Isomorphism Theorem

Given an infinite Markovian trajectory X_1^∞ and $t \geq 1$, let $\{N_i^X(t), \forall i \in [n]\}$ be the counting measure of states $[n]$ appearing in the subtrajectory X_1^t up to time t, and we denote the empirical distribution induced by the trajectory as $\hat{\pi}(t) = (N_1(t)/t, ..., N_n(t)/t)$. We define the random cover time as $\tau_{\text{cov}}^X \triangleq \inf\{t : \forall i \in [n], N_i^X(t) > 0\}$, the first time to have visited every state. For clearer illustration, we omit the superscript X in $N_i^X(t)$ and τ_{cov}^X in the rest of the paper when it does not incur ambiguity. The expectation of τ_{cov} given a fixed initial state i_0 is $\mathbb{E}[\tau_{\text{cov}}|X_0 = i_0]$ and the expected cover time is the maximum over the initial state, $t_{\text{cov}} \triangleq \max_{i_0 \in [n]} \mathbb{E}[\tau_{\text{cov}}|X_0 = i_0]$.

The *random hitting time* is the first time when a certain state gets hit by the random walk. Specifically, for some $j \in [n]$ the random hitting time is $\tau_{\text{hit}}(j) \triangleq \inf\{t : N_j(t) > 0\}$. The hitting time is then defined as $t_{\text{hit}} = \max_{i_0, j \in [n]} \mathbb{E}[\tau_{\text{hit}}(j)|X_0 = i_0]$.

For any reversible Markov chain, denote $\Pi = Diag(\pi)$, then $A = \Pi M$ is a symmetric matrix due to reversibility. Specially, associate with A an undirected graph $G = (V, E, w)$ with $|V| = n$ and edge weights $w_{ij} \propto \pi_i p_{ij}$, which follows the matrix A. Then the Markov chain is connected to the random walk over this undirected weighted graph via $p_{ij} = w_{ij} / \sum_j w_{ij}$. We denote $c_i = \sum_j w_{ij} \propto \pi_i$ as the capacity of node i and consider each edge with weight w_{ij} as a wire with capacitance w_{ij}, or equivalently, a wire with resistance $1/w_{ij}$. Then we can also think of this graph as a resistance network. Given any two node i, j, we denote the *effective resistance* between i, j over this network as r_{ij}. Specially, we denote $c = \sum_i c_i$ as the total capacity.

The continuous-time Markov chain can be constructed from a discrete Markov chain by setting an exponential clock $\tau_{\text{exp}} \sim \text{Exp}(1)$ to determine the time interval between jumps. Fix the starting state as $i_0 \in [n]$, the *local time* for state $i \in [n]$ and time t is

$$L_t^i \triangleq \frac{1}{c_i} \int_0^t \mathbf{1}_{\{X_s = i\}} ds.$$

The *inverse local time* for state i is

$$\tau_{\text{inv}}(t) \triangleq \inf\{s : L_s^i > t\},$$

where the dependence on i is hidden in the notation and will be clear from the context. Following Ding et al. (2011); Ding (2014), we will analyze the (k-) cover time via the local time process $\left\{L_{\tau_{\text{inv}}(t)}^i : i \in [n]\right\}$.

We now recall the generalized second Ray-Knight isomorphism theorem of Eisenbaum et al. (2000) (see also (Marcus and Rosen, 2006, Theorem 8.2.2)).

Theorem 2 (Generalized Second Ray-Knight isomorphism theorem) *Fix some state $i_0 \in [n]$ and denote $T_0 \triangleq \tau_{\text{hit}}(i_0)$. We let*

$$\Gamma_{i_0}(i, j) = \mathbb{E}[L_{T_0}^j | X_0 = i] = \frac{1}{2}(r_{i_0 i} + r_{i_0 j} - r_{ij})$$

and let $\eta = \{\eta_i : i \in [n]\}$ be a mean zero Gaussian process with covariance $\Gamma_{i_0}(i, j)$. Let P_{i_0} and P_η be the measure on the process $\left\{L_{\tau_{\text{inv}}(t)}^i\right\}$ and $\{\eta_x\}$, respectively. Then under the measure $P_{i_0} \times P_\eta$, for any $t > 0$, we have the following equality in distribution:

$$\left\{L_{\tau_{\text{inv}}(t)}^i + \frac{1}{2}\eta_i^2 : i \in [n]\right\} \stackrel{d.}{=} \left\{\frac{1}{2}(\eta_i' + \sqrt{2t})^2 : i \in [n]\right\}.$$

This powerful isomorphism theorem was used by Ding et al. (2011) to prove the "blanket time conjecture" of Winkler and Zuckerman (1996). And the Gaussian process described above is called the *Gaussian free field* in the literature. We cite the main theorem of Ding et al. (2011) for future reference.

Theorem 3 (Constant-factor approximation of cover time) *For the random walk on reversible Markov chains, fix some $i_0 \in [n]$ as starting state, and let $\eta = \{\eta_i : i \in [n]\}$ be the Gaussian process described in Theorem 2. Then we have*

$$t_{\text{cov}} \asymp c \left(\mathbb{E} \max_i \eta_i \right)^2 .$$

The k-cover time naturally generalizes the cover time, and underpins our arguments for sample complexity bounds. Roughly speaking, it measures the expected length of the Markovian trajectory to ensure covering each state k times.

Definition 4 (k-cover time) *For any $k \in \mathbf{N}^+$, the random k-cover time $\tau_{\text{cov}}^{(k)}$ is the first time when every state in $[n]$ has been visited k times, i.e., $\tau_{\text{cov}}^{(k)} \triangleq \inf\{t : \forall i \in [n], N_i(t) \geq k\}$. And the k-cover time is $t_{\text{cov}}^{(k)} \triangleq \max_{i_0 \in [n]} \mathbb{E}[\tau_{\text{cov}}^{(k)} | X_0 = i_0]$.*

Note that the k-cover time coincides with the cover time when $k = 1$. And we refer the readers to Levin and Peres (2017); Aldous and Fill (1995) for a wonderful exposition of techniques and results on Markov chains.

The rest of this paper is structured as follows. In Section 3, we connect Markov chain learning/testing to k-cover time. In Section 4, we bound the k-cover time of reversible chains via the isomorphism theorem, and discuss its implications on testing and learning. In Section 5, we bound the k-cover time of irreducible chains, discuss its consequences and end with several open problems.

3. Learning and Testing Markov Chains via k-cover Time

In this section, we will see how k-cover time is closely related to Markov chain learning and testing problems. In the following, we argue that if $k(n, \epsilon, \delta)$ i.i.d. samples are enough to (ϵ, δ)-learn/test n-state discrete distributions under total variation distance, then $t_{\text{cov}}^{k(n,\epsilon,O(\delta/n))}$ samples are sufficient to learn/test the Markov chain under infinity matrix norm. We first prove the following simple lemma.

Lemma 5 (Exponential decay lemma) *For random walk on irreducible chains, for any $k, m \in \mathbf{N}^+$, and any initial distribution $\mathbf{q}$, we have $\mathbb{P}(\tau_{\text{cov}}^{(k)} \geq emt_{\text{cov}}^{(k)}) \leq e^{-m}$.*

Proof Consider $\tau_{\text{cov}}^{(k)}$ with any fixed starting state $X_0 \sim \mathbf{q}$, we have by Markov's inequality and linearity of expectation that

$$\mathbb{P}(\tau_{\text{cov}}^{(k)} \geq et_{\text{cov}}^{(k)}) \leq \mathbb{P}(\tau_{\text{cov}}^{(k)} \geq e\, \mathbb{E}[\tau_{\text{cov}}^{(k)} | X_0]) \leq 1/e. \tag{1}$$

Note that this inequality holds for any initial distribution of starting state $\mathbf{q}$. We then bound $\mathbb{P}(\tau_{\text{cov}}^{(k)} \leq emt_{\text{cov}}^{(k)}) \geq e^{-m}$ by induction.

First, we consider the first two sub-trajectories of the Markov chain, each of length $l \triangleq et_{\text{cov}}^{(k)}$, i.e., the chain X_1^l and X_{l+1}^{2l}. Denote the event $E_1 \triangleq \{X_1^l \text{ covers the state space } k \text{ times}\}$, and $E_2 \triangleq \{X_{l+1}^{2l} \text{ covers the state space } k \text{ times}\}$. Suppose X_0 is drawn from $\mathbf{q} \in \Delta^{n-1}$, then according to Eq. (1), we have $\mathbb{P}(E_1^c) = \mathbb{P}(\tau_{\text{cov}}^{(k)} \geq et_{\text{cov}}^{(k)}) \leq 1/e$. Denote the distribution of X_l conditioned on E_1^c as $\mathbf{q}'$, and $\tau_{\text{cov}}^{(k)'}$ as the k-cover time of X_{l+1}^{∞}, then we have $\mathbb{P}(E_2^c|E_1^c) = \mathbb{P}(\tau_{\text{cov}}^{(k)'} \geq et_{\text{cov}}^{(k)}|\tau_{\text{cov}}^{(k)} \geq et_{\text{cov}}^{(k)}) = \mathbb{P}(\tau_{\text{cov}}^{(k)'} \geq et_{\text{cov}}^{(k)}|X_l \sim \mathbf{q}') \leq 1/e$. Here we used the fact that E_1 is determined by X_1^l; while due to Markovian property, E_2 do not depend on X_1^{l-1}.

The above reasoning gives $\mathbb{P}(E_1^c \cap E_2^c) = \mathbb{P}(E_1^c)\,\mathbb{P}(E_2^c|E_1^c) \leq e^{-2}$. Similarly, we can deduce that $\mathbb{P}(\cap_{i\in[m]} E_i^c) \leq e^{-m}$. But the event $E \triangleq \{X_1^{ml} \text{ covers the state space } k \text{ times}\}$ includes the event $\cup_{i\in[m]} E_i$, thus $\mathbb{P}(E^c) \leq \mathbb{P}(\cap_{i\in[m]} E_i^c) \leq e^{-m}$. This proves the lemma. ■

3.1. Learning Markov Chains

Given *any* (ϵ, δ)-learner $\mathcal{L}(Y_1^m, n)$ for discrete distributions that outputs $\hat{\mathbf{p}}$ with i.i.d. samples Y_1^m from $\mathbf{p} \in \Delta^{n-1}$, we consider the following learning algorithm for Markov chains. Here $k(n, \epsilon, \delta)$ is the sample complexity of (ϵ, δ)-learn a discrete distribution using i.i.d. samples.

Input: a Markovian trajectory X_1^m, parameters n, ϵ, δ
Output: a candidate Markov chain $\hat{M}$
for $i \leftarrow 1, 2, ..., n$ **do**
  if $N_i^X(m) \leq k(n, \epsilon, \delta/2n)$ **then**
    $\hat{\mathbf{p}}_i \leftarrow \frac{1}{n}\mathbf{1}$
  else
    Let $Y_{i,1}, Y_{i,2}, ..., Y_{i,k(n,\epsilon,\delta/2n)}$ be the first $k(n, \epsilon, \delta/2n)$ succeeding states of state i in X_1^m
    $\hat{\mathbf{p}}_i \leftarrow \mathcal{L}((Y_{i,1}, ..., Y_{i,k(n,\epsilon,\delta/2n)}), n)$
  end
end
return $\hat{M} \leftarrow (\hat{\mathbf{p}}_1, ..., \hat{\mathbf{p}}_n)$

Algorithm 1: LEARNCHAIN

Then we have the following lemma about the sample complexity of learning Markov chains.

Lemma 6 (k-cover time and learning Markov chain) *If we have a (ϵ, δ)-learner for n-state distribution with sample complexity $k(n, \epsilon, \delta)$, then we can (ϵ, δ)-learning the chain M using $O_\delta(t_{\text{cov}}^{k(n,\epsilon,\delta/2n)})$ samples. Here O_δ hides logarithmic factors in δ.*

Proof Since we have $\mathbb{P}(\tau_{\text{cov}}^{(k)} \geq emt_{\text{cov}}^{(k)}) \leq e^{-m}$ according to Theorem 5, then by taking $m = \ln\frac{2}{\delta}$, we have $\mathbb{P}(\tau_{\text{cov}}^{(k)} \geq et_{\text{cov}}^{(k)} \ln\frac{2}{\delta}) \leq \frac{\delta}{2}$. Thus, for a length $l = et_{\text{cov}}^{k(n,\epsilon,\delta/2n)} \ln\frac{2}{\delta}$ trajectory, we will have $k(n, \epsilon, \delta/2n)$ samples for each states in $[n]$ with probability $\geq 1 - \delta/2$. We consider the infinite chain X_1^∞, and define the event $E = \{N_i(l) \geq k(n, \epsilon, \delta/2n)\}$, $E_i = \{$ first k samples for state i from X_1^∞ yields $d_{\text{TV}}(\hat{\mathbf{p}}_i, \mathbf{p}_i) \leq \epsilon\}$. Then $\mathbb{P}(E) \geq 1 - \delta/2$, and $\mathbb{P}(E^c) \leq \delta/2$; also we have $\mathbb{P}(E_i) \geq 1 - \delta/2n$, and $\mathbb{P}(E_i^c) \leq \delta/2n$, due to the Markov property and the guarantee of the discrete distribution learner $\mathcal{L}$.

This gives that $\mathbb{P}(E \cap E_1 ... \cap E_n) = 1 - \mathbb{P}(E^c \cup E_1^c ... \cup E_n^c)$. But by union bound $\mathbb{P}(E^c \cup E_1^c ... \cup E_n^c) \leq \mathbb{P}(E^c) + \sum_{i=1}^n \mathbb{P}(E_i^c) \leq \delta$. And $E \cup E_1 ... \cup E_n$ implies that we have for all

$i \in [n]$, $d_{\text{TV}}(\hat{\mathbf{p}}_i, \mathbf{p}_i) \le \epsilon$, which guarantees $\|\hat{M} - M\|_\infty = \max_{i\in[n]} d_{\text{TV}}(\hat{\mathbf{p}}_i, \mathbf{p}_i) \le \epsilon$. Thus, with probability $\ge 1-\delta$, we will have both $\tau_{\text{cov}}^{(k)} \le et_{\text{cov}}^{(k)} \ln \frac{2}{\delta}$ and $\|\hat{M} - M\|_\infty \le \epsilon$. Therefore, we can (ϵ, δ)-learn the chain using $O_\delta(t_{\text{cov}}^{k(n,\epsilon,\delta/2n)})$ samples. ■

3.2. Identity Testing of Markov Chains

We now consider the task of identity testing of Markov chains. Given *any* (ϵ, δ)-identity-tester $\mathcal{T}(Y_1^m, n, \mathbf{p})$ for discrete distributions that outputs "Yes" if $\mathbf{p} = \mathbf{p}'$ and "No" if $d_{\text{TV}}(\mathbf{p}, \mathbf{p}') \ge \epsilon$, we consider the following identity testing algorithm for Markov chains. Here $k(n, \epsilon, \delta)$ is the sample complexity of (ϵ, δ)-identity-test a discrete distribution using i.i.d. samples.

Input: a Markovian trajectory X_1^m, parameters n, ϵ, δ, a reference chain M
Output: "Yes" if $M = M'$, "No" if $\|M - M'\|_\infty \ge \epsilon$
for $i \leftarrow 1, 2, ..., n$ **do**
  if $N_i^X(m) \le k(n, \epsilon, \delta/2n)$ **then**
    return *"No"*
  else
    Let $Y_{i,1}, Y_{i,2}, ..., Y_{i,k(n,\epsilon,\delta/2n)}$ be the first $k(n, \epsilon, \delta/2n)$ succeeding states of state i in X_1^m **if** $\mathcal{T}((Y_{i,1}, ..., Y_{i,k(n,\epsilon,\delta/2n)}), n, \mathbf{p}_i) =$ *"No"* **then**
      return *"No"*
    end
  end
end
return *"Yes"*

Algorithm 2: IDTESTCHAIN

Similarly, we have the following lemma about the sample complexity of identity-testing Markov chains.

Lemma 7 (k-cover time and identity-testing Markov chain) *If we have a (ϵ, δ)-identity-tester for n-state distribution with sample complexity $k(n, \epsilon, \delta)$, then we can (ϵ, δ)-identity-testing the chain M against unknown chain M' using $O_\delta(t_{\text{cov}}^{k(n,\epsilon,\delta/2n)}(M))$ samples. Here O_δ hides logarithmic factors in δ, and we use $t_{\text{cov}}^{(k)}(M)$ to specify the k-cover time of M instead of M'.*

Proof We consider two cases (i) $M = M'$ and (ii) $\|M - M'\|_\infty \ge \epsilon$ as the following. Similarly, we consider the infinite chain X_1^∞, and denote $E = \{N_i(l) \ge k(n, \epsilon, \delta/2n)\}$, $E_i = \{$ the first k samples for state i from X_1^∞ yields "No" during the test $\}$.
Case 1. $M = M'$.

Due to Theorem 5, for a length $l = et_{\text{cov}}^{k(n,\epsilon,\delta/2n)} \ln \frac{2}{\delta}$ trajectory, we will have $k(n, \epsilon, \delta/2n)$ samples for each state with probability $\ge 1 - \delta/2$, thus $\mathbb{P}(E) \ge 1 - \delta/2$. Moreover, by Markov property and the guarantee of the learner, the event E_i happens with probability $\mathbb{P}(E_i) \le \delta/2n$ for any $i \in [n]$. Thus by a union bound, error events happen with probability $\mathbb{P}(E^c \cup E_1 ... \cup E_n) \le \delta$. And with probability $\ge 1 - \delta$, the identity tester will answer "Yes".
Case 2. $\|M - M'\|_\infty \ge \epsilon$.

The only case it makes fault by answering "Yes" is when it do not pass Line 2 and Line 6 for all states, which means it will have enough samples for testing each state, and the i.i.d. tester $\mathcal{T}$ answers "Yes" for all sub-tests $\{\mathbf{p}_i, \forall i \in [n]\}$. Since $\|M - M'\|_\infty \geq \epsilon$ implies there exists $i_* \in [n]$ such that $d_{\text{TV}}(\mathbf{p}_{i_*}, \mathbf{p}'_{i_*}) \geq \epsilon$, and this guarantees that the sub-test for i_* will return "No" with probability $\mathbb{P}(E_{i^*}) \geq 1 - \delta/2n$. Thus the probability of the whole process answering "Yes" is $\mathbb{P}(E \cap E_1^c ... \cap E_n^c) \leq \mathbb{P}(E_{i^*}^c) \leq \delta/2n$.

To sum up, for both cases, the identity tester will give the correct answer with probability $\geq 1 - \delta$. This proves the lemma. ■

3.3. Closeness Testing of Markov Chains

We now considering the task of closeness testing of Markov chains. Given *any* (ϵ, δ)-closeness-tester $\mathcal{T}(Y_1^m, Y_1^{m'}, n)$ for discrete distributions that outputs "Yes" if $\mathbf{p} = \mathbf{p}'$ and "No" if $d_{\text{TV}}(\mathbf{p}, \mathbf{p}') \geq \epsilon$, we consider the following identity testing algorithm for Markov chains, where $k(n, \epsilon, \delta)$ is the sample complexity of (ϵ, δ)-closeness-test a discrete distribution using i.i.d. samples.

```
Input: two Markovian trajectories X_1^m, X_1^{m'}, parameters n, ε, δ
Output: "Yes" if M = M', "No" if ||M − M'||_∞ ≥ ε
for i ← 1, 2, ..., n do
    if N_i^X(m) ≤ k(n, ε, δ/2n) or N_i^{X'}(m) ≤ k(n, ε, δ/2n) then
        return "No"
    else
        Let Y_{i,1}, Y_{i,2}, ..., Y_{i,k(n,ε,δ/2n)} be the first k(n, ε, δ/2n) succeeding states of state i in X_1^m
        Let Y'_{i,1}, Y'_{i,2}, ..., Y'_{i,k(n,ε,δ/2n)} be the first k(n, ε, δ/2n) succeeding states of state i in X_1^{m'}
        if T((Y_{i,1}, Y_{i,2}, ..., Y_{i,k(n,ε,δ/2n)}), (Y'_{i,1}, ..., Y'_{i,k(n,ε,δ/2n)}), n) = "No" then
            return "No"
        end
    end
end
return "Yes"
```

Algorithm 3: CLOSETESTCHAIN

We then have the following lemma connecting k-cover time to the sample complexity of closeness-testing Markov chains.

Lemma 8 (k-cover time and closeness-testing Markov chain) *If we have a (ϵ, δ)-closeness-tester for n-state distribution with sample complexity $k(n, \epsilon, \delta)$, then we can (ϵ, δ)-closeness-testing the unknown chains M, M' using $O_\delta(\min\{t_{\text{cov}}^{k(n,\epsilon,\delta/4n)}(M), t_{\text{cov}}^{k(n,\epsilon,\delta/4n)}(M')\})$ samples. Here O_δ hides logarithmic factors in δ.*

Proof Consider the cases (i) $M = M'$ and (ii) $\|M - M'\|_\infty \geq \epsilon$ as the following. Consider the infinite chain X_1^∞, and denote $E_X = \{N_i^X(l) \geq k(n, \epsilon, \delta/4n)\}$, $E_Y = \{N_i^Y(l) \geq k(n, \epsilon, \delta/4n)\}$, $E_i = \{$ the first k samples for state i from X_1^∞ yields "No" during the test $\}$.

Case 1. $M = M'$.

Due to Theorem 5, for a length $l = e t_{\text{cov}}^{k(n,\epsilon,\delta/4n)} \ln \frac{2}{\delta}$ trajectory, we will have $k(n, \epsilon, \delta/4n)$ samples for each state with probability $\mathbb{P}(E_X) \geq 1 - \delta/4$ and $\mathbb{P}(E_Y) \geq 1 - \delta/4$. By a union bound over the two chains, the probability of passing the condition in Line 2 of Algorithm 2 is $\geq 1-\delta/2$. Then we have $\mathbb{P}(E_i) \leq \delta/4n$, and error probability $\mathbb{P}(E_X^c \cup E_Y^c \cup E_1 ... \cup E_n) \leq 3\delta/4$. Thus with probability $\geq 1-\delta$ the identity tester will answer "Yes".

Case 2. $\|M - M'\|_\infty \geq \epsilon$.

The only case it answers "Yes" is when it do not pass Line 2 and Line 7 in Algorithm 3 for all states, which means it will have enough samples for testing each state, and the i.i.d. tester $\mathcal{T}$ answers "Yes" for all sub-tests. Then essentially the same argument in Theorem 7 will give that $\mathbb{P}(E_{i^*}) \geq 1 - \delta/4n$ for some i^*. Therefore, the probability of answering "Yes" is $\mathbb{P}(E_X \cap E_Y \cap E_1^c ... \cap E_n^c) \leq \mathbb{P}(E_{i^*}^c) \leq \delta/4n$. This proves the lemma. ■

3.4. Tolerant Testing and More

We now considering the task of tolerant identity/closeness testing of Markov chains. Given *any* $(\epsilon_1, \epsilon_2, \delta)$-tolerant-identity-tester $\mathcal{T}(X_1^m, n, \mathbf{p})$ for discrete distributions that outputs "Yes" if $d_{\text{TV}}(\mathbf{p}, \mathbf{p}') \leq \epsilon_1$ and "No" if $d_{\text{TV}}(\mathbf{p}, \mathbf{p}') \geq \epsilon_2$, we can construct similar tolerant tester for Markov chains as above. We have the following propositions for tolerant testing problems.

Lemma 9 (k-cover time and tolerant-identity-testing Markov chain) *If we have a $(\epsilon_1, \epsilon_2, \delta)$-tolerant-identity-tester for n-state distribution with sample complexity $k(n, \epsilon_1, \epsilon_2, \delta)$, then we can $(\epsilon_1, \epsilon_2, \delta)$-tolerant-identity-testing M against the unknown chains M using*

$$O_\delta(\max\{t_{\text{cov}}^{k(n,\epsilon_1,\epsilon_2,\delta/2n)}(M), t_{\text{cov}}^{k(n,\epsilon_1,\epsilon_2,\delta/2n)}(M')\})$$

samples.

Lemma 10 (k-cover Time and Tolerant-closeness-testing Markov Chain) *If we have a $(\epsilon_1, \epsilon_2, \delta)$-tolerant-closeness-tester for n-state distribution with sample complexity $k(n, \epsilon_1, \epsilon_2, \delta)$, then we can $(\epsilon_1, \epsilon_2, \delta)$-tolerant-closeness-testing the unknown chains M, M' using*

$$O_\delta(\max\{t_{\text{cov}}^{k(n,\epsilon_1,\epsilon_2,\delta/4n)}(M), t_{\text{cov}}^{k(n,\epsilon_1,\epsilon_2,\delta/4n)}(M')\})$$

samples.

Besides these, we also have the problem of testing with respect to uniform distributions. We have the following problem for the Markov chain scenario. Given a trajectory X_1^m from M', can we test whether it comes from uniform distribution $M = \frac{1}{n}\mathbf{1}\mathbf{1}^T$, or it comes from M' such that $\|M' - M\|_\infty \geq \epsilon$. Then we have the following propositions.

Lemma 11 (k-cover Time and Uniform-testing Markov Chain) *If we have a (ϵ, δ)-uniform-tester for n-state distribution with sample complexity $k(n, \epsilon, \delta)$, then we can (ϵ, δ)-uniform-testing against unknown chain M' using $O_\delta(t_{\text{cov}}^{k(n,\epsilon,\delta/2n)}(M))$ samples, where $M = \frac{1}{n}\mathbf{1}\mathbf{1}^T$. We remark that it does not depend on the k-cover time of the unknown chain.*

Lemma 12 (k-cover Time and Tolerant-uniform-testing Markov Chain) *If we have a $(\epsilon_1, \epsilon_2, \delta)$-tolerant-uniform-tester for n-state distribution with sample complexity $k(n, \epsilon_1, \epsilon_2, \delta)$, then we can $(\epsilon_1, \epsilon_2, \delta)$-tolerant-uniform-testing against the unknown chains M' using $O_\delta(t_{\text{cov}}^{k(n,\epsilon_1,\epsilon_2,\delta/2n)}(M'))$ samples.*

As the above arguments show, k-cover time establishes a universal connection between the testing and learning problems of Markov chains and discrete distributions. In a sense, the Markov chain learning/testing problems can be reduced to those over discrete distributions via k-cover time. Thus an interesting question would be to bound the k-cover time, in terms of basic quantities like n, π_* and t_{cov} associated with a Markov chain. In the next section, we will prove that $t_{\text{cov}}^{(k)} = \Theta(t_{\text{cov}} + k/\pi_*)$ for reversible chains and $t_{\text{cov}}^{(k)} = \tilde{\Theta}(t_{\text{cov}} + k/\pi_*)$ for irreducible chains. These bounds on k-cover time then gives nice sample complexity bounds on learning/testing Markov chain problems in an unified version.

4. The k-cover Time of Reversible Chains

In this section, we focus on bounding the k-cover time of reversible Markov chains with respect to the basic quantities n, π_* and t_{cov}. First, we prove an universal lower bound of k-cover time that applies to all irreducible Markov chains. Then we prove a tight upper bound on $t_{\text{cov}}^{(k)}$ for reversible Markov chains.

4.1. Lower Bound for General Irreducible Chains

We have the following lower bound on $t_{\text{cov}}^{(k)}$ for all irreducible Markov chains. To prove the lemma, we will use the connection to *return time*. For some state $i \in [n]$, the return time $\tau_{\text{ret}}(i)$ is the first time a Markov chain starting at i returns to i. And the expected return time is $t_{\text{ret}}(i) = \mathbb{E}[\tau_{\text{ret}}|X_0 = i]$. It is standard result that $t_{\text{ret}}(i) = 1/\pi_i$.

Lemma 13 (lower bound on k-cover time) *For any irreducible Markov chain with minimum stationary probability π_* and cover time t_{cov}, we have $t_{\text{cov}}^{(k)} = \Omega(k/\pi_* + t_{\text{cov}})$.*

Proof Clearly we have $t_{\text{cov}}^{(k)} \geq t_{\text{cov}}$ for all $k \geq 1$. We will show $t_{\text{cov}} \geq (k-1)/\pi_*$, which proves the lemma. Denote $i_* = \arg\min_{i\in[n]} \pi_i$ and the pth time of hitting state i_* as $\tau_{\text{hit}}^{(p)}(i_*)$, then it's clear that $\tau_{\text{cov}}^{(k)} \geq \tau_{\text{hit}}^{(k)}(i_*)$ for any chain. Thus $\mathbb{E}[\tau_{\text{cov}}^{(k)}|X_0 = i_0] \geq \mathbb{E}[\tau_{\text{hit}}^{(k)}(i_*)|X_0 = i_0]$. Note that $\tau_{\text{hit}}^{(k)}(i_*) = \tau_{\text{hit}}(i_*) + \sum_{j=2}^{k} \left(\tau_{\text{hit}}^{(j)}(i_*) - \tau_{\text{hit}}^{(j-1)}(i_*)\right)$. Then we have

$$\begin{aligned}
\mathbb{E}[\tau_{\text{hit}}^{(k)}(i_*)|X_0 = i_0] &= \mathbb{E}\left[\tau_{\text{hit}}(i_*) + \sum_{j=2}^{k} \left(\tau_{\text{hit}}^{(j)}(i_*) - \tau_{\text{hit}}^{(j-1)}(i_*)\right)|X_0 = i_0\right] \\
&= \mathbb{E}[\tau_{\text{hit}}(i_*)|X_0 = i_0] + \sum_{j=2}^{k} \mathbb{E}\left[\tau_{\text{hit}}^{(j)}(i_*) - \tau_{\text{hit}}^{(j-1)}(i_*)|X_0 = i_0\right] \\
&\geq \sum_{j=2}^{k} \mathbb{E}\left[\tau_{\text{hit}}^{(j)}(i_*) - \tau_{\text{hit}}^{(j-1)}(i_*)|X_0 = i_0\right]
\end{aligned} \tag{2}$$

Note that due to the Markov property, $\tau_{\text{hit}}^{(j)}(i_*) - \tau_{\text{hit}}^{(j-1)}(i_*)$ in fact has the same distribution as $\tau_{\text{ret}}(i_*)$. This certifies that $\mathbb{E}\left[\tau_{\text{hit}}^{(j)}(i_*) - \tau_{\text{hit}}^{(j-1)}(i_*)|X_0 = i_0\right] = (k-1)t_{\text{ret}}(i_*) = (k-1)/\pi_*$. Thus, we have $t_{\text{cov}}^{(k)} \geq (k-1)/\pi_*$ and the lemma is proved. ■

4.2. Upper Bound for Reversible Chains

Now we prove a matching upper bound for reversible Markov chains using the Ray–Knight's isomorphism theorem. The following lemma on the concentration of the supremum of a Gaussian process is useful Ding et al. (2011).

Lemma 14 (Gaussian supremum lemma) *Consider a Gaussian process $\{\eta_i : i \in [n]\}$ and define $\sigma = \sup_{i\in[n]} \sqrt{\mathbb{E}\,\eta_i^2}$. Then for $\alpha > 0$, we have*

$$\mathbb{P}\left(\left|\sup_{i\in[n]} \eta_i - \mathbb{E}\sup_{i\in[n]} \eta_i\right| > \alpha\right) \leq 2\exp(-\alpha^2/2\sigma^2).$$

Also, we will use the concentration of the inverse local time (Ding, 2014, Lemma 2.1).

Lemma 15 (Inverse local time lemma) *Let X be a continuous time random walk on an electrical network, and denote $c = \sum_{i,j\in[n]} w_{ij}$ be the total conductance. Fixing any state $i_0 \in [n]$, let $R \triangleq \max_{i,j\in[n]} \mathbb{E}(\eta_i - \eta_j)^2$ and $\tau_{\text{inv}}(t)$ be the inverse local time for i_0. Then for any $t \geq 0$ and $\lambda \geq 1$,*

$$\mathbb{P}\left(\left|\tau_{\text{inv}}(t) - c\cdot t\right| \geq \frac{1}{2}(\sqrt{\lambda R t} + \lambda R)\cdot c\right) \leq 6\exp(-\lambda/16).$$

Armed with this lemma, then we can prove an upper bound on the k-cover time of reversible Markov chains for the continuous-time scenario as the following.

Theorem 16 (k-cover time of continuous-time reversible chains) *For continuous-time random walk on reversible chains, we have $t_{\text{cov}}^{(k)} = O(k/\pi_* + t_{\text{cov}})$.*

Proof We fix any $i_0 \in [n]$, and let $\tau_{\text{inv}}(t)$ be the corresponding inverse local time for i_0. Note by $\tau_{\text{inv}}(t)$, with high probability, we should have accumulated $\Omega(t)$ local time at each node. To show this, we denote $\Lambda = \mathbb{E}\sup_i \eta_i$ and for some small constant $\delta \in (0,1)$, consider the event $E = \{\inf_i L^i_{\tau_{\text{inv}}(t)} \leq \delta t\}$. Note that by the isomorphism theorem (Theorem 2),

$$\left\{L^i_{\tau_{\text{inv}}(t)} + \frac{1}{2}\eta_i^2 : i \in [n]\right\} \overset{d.}{=} \left\{\frac{1}{2}(\eta_i' + \sqrt{2t})^2 : i \in [n]\right\}.$$

Thus we have

$$\mathbb{P}\left(\inf_i L^i_{\tau_{\text{inv}}(t)} + \frac{1}{2}\eta_i^2 \leq (1+\delta)t/2\right) = \mathbb{P}\left(\inf_i \frac{1}{2}(\eta_i' + \sqrt{2t})^2 \leq (1+\delta)t/2\right) .$$

And we also have

$$\mathbb{P}\left(\inf_i L^i_{\tau_{\text{inv}}(t)} + \frac{1}{2}\eta_i^2 \leq (1+\delta)t/2\right) \geq \mathbb{P}\left(\inf_i L^i_{\tau_{\text{inv}}(t)} + \sup_i \frac{1}{2}\eta_i^2 \leq (1+\delta)t/2\right) .$$

Moreover, suppose $\inf_i L^i_{\tau_{\text{inv}}(t)} \leq \delta t$ and $\inf_i L^i_{\tau_{\text{inv}}(t)} + \sup_i \frac{1}{2}\eta_i^2 \geq (1+\delta)t/2$, then we must have

$$\sup_i \frac{1}{2}\eta_i^2 \geq (1+\delta)t/2 - \inf_i L^i_{\tau_{\text{inv}}(t)} \geq (1-\delta)t/2.$$

This shows that

$$\mathbb{P}(E) \leq \mathbb{P}(\sup_i \frac{1}{2}\eta_i^2 \geq (1-\delta)t/2 \text{ or } \inf_i L^i_{\tau_{\text{inv}}(t)} + \sup_i \frac{1}{2}\eta_i^2 \leq (1+\delta)t/2).$$

By union bound and previous inequalities we have

$$\begin{aligned}\mathbb{P}(E) &\leq \mathbb{P}(\sup_i \frac{1}{2}\eta_i^2 \geq (1-\delta)t/2) + \mathbb{P}(\inf_i \frac{1}{2}(\eta_i' + \sqrt{2t})^2 \leq (1+\delta)t/2) \\ &\leq \mathbb{P}(\sup_i |\eta_i| \geq \sqrt{(1-\delta)t}) + \mathbb{P}(\inf_i \eta_i' \leq \sqrt{(1+\delta)t} - \sqrt{2t})\end{aligned} \tag{3}$$

Here we used the fact that $\inf_i |\eta_i' + \sqrt{2t}| \geq \inf_i \eta_i' + \sqrt{2t}$. Note that by symmetry of centered Gaussian process, we have

$$\mathbb{P}(\inf_i \eta_i' \leq \sqrt{(1+\delta)t} - \sqrt{2t}) = \mathbb{P}(\sup_i \eta_i' \geq \sqrt{2t} - \sqrt{(1+\delta)t})$$

and

$$\begin{aligned}\mathbb{P}(\sup_i |\eta_i| \geq \sqrt{(1-\delta)t}) &= \mathbb{P}(\sup_i \eta_i \geq \sqrt{(1-\delta)t} \text{ or } \inf_i \eta_i \leq -\sqrt{(1-\delta)t}) \\ &\leq \mathbb{P}(\sup_i \eta_i \geq \sqrt{(1-\delta)t}) + \mathbb{P}(\inf_i \eta_i \leq -\sqrt{(1-\delta)t}) \\ &= 2\,\mathbb{P}(\sup_i \eta_i \geq \sqrt{(1-\delta)t})\end{aligned} \tag{4}$$

Now by concentration of the supremum of a Gaussian process, we deduce that for $t \geq \Lambda^2/(1-\delta)$,

$$\begin{aligned}\mathbb{P}(\sup_i \eta_i \geq \sqrt{(1-\delta)t}) &= \mathbb{P}(\sup_i \eta_i - \Lambda \geq \sqrt{(1-\delta)t} - \Lambda) \\ &\leq \mathbb{P}(|\sup_i \eta_i - \Lambda| \geq \sqrt{(1-\delta)t} - \Lambda) \\ &\leq 2\exp(-(\sqrt{(1-\delta)t} - \Lambda)^2/2\sigma^2).\end{aligned} \tag{5}$$

Similarly, we have for $t \geq \Lambda^2/(\sqrt{2} - \sqrt{1+\delta})^2$,

$$\begin{aligned}\mathbb{P}(\sup_i \eta_i' \geq \sqrt{2t} - \sqrt{(1+\delta)t}) &= \mathbb{P}(\sup_i \eta_i' - \Lambda \geq \sqrt{2t} - \sqrt{(1+\delta)t} - \Lambda) \\ &\leq 2\exp(-(\sqrt{2t} - \sqrt{(1+\delta)t} - \Lambda)^2/2\sigma^2).\end{aligned} \tag{6}$$

To this end, we have shown that for $\delta = 1/2$, and $t \geq 8\Lambda^2/(2-\sqrt{3})^2 \simeq 228.6\Lambda^2$,

$$\begin{aligned}\mathbb{P}(E) &\leq 4\exp(-(\sqrt{t} - \sqrt{2}\Lambda)^2/4\sigma^2) + 2\exp(-((2-\sqrt{3})\sqrt{t} - \sqrt{2}\Lambda)^2/4\sigma^2) \\ &\leq 6\exp(-((2-\sqrt{3})\sqrt{t} - \sqrt{2}\Lambda)^2/4\sigma^2) \\ &\leq 6\exp(-(2-\sqrt{3})^2 t/16\sigma^2) \\ &\leq 6\exp(-t/450\sigma^2)\end{aligned} \tag{7}$$

Finally, we have shown that $\mathbb{P}(\min_i L^i_{\tau_{\text{inv}}(t)} \leq t/2) \leq 6\exp(-t/450\sigma^2)$ for $t \geq 230\Lambda^2$. Now we will use the concentration of the inverse local time.

$$\mathbb{P}\left(\left|\tau_{\text{inv}}(t) - c \cdot t\right| \geq (\sqrt{\lambda R t} + 2\lambda R) \cdot c\right) \leq 6\exp(-\lambda/4)$$

Note $R = \max_{i,j\in[n]} \mathbb{E}(\eta_i - \eta_j)^2$, hence

$$\sigma^2 = \max_{j\in[n]} \mathbb{E}(\eta_{i_0} - \eta_j)^2 \leq R \leq \max_{i,j\in[n]} 2\,\mathbb{E}(\eta_i^2 + \eta_j^2) = 4\sigma^2.$$

Specially, we have

$$\mathbb{P}\left(\tau_{\text{inv}}(t) \geq ct + c(2\sigma\sqrt{\lambda t} + 8\lambda\sigma^2)\right) \leq 6\exp(-\lambda/4).$$

Taking $\lambda = t/100\sigma^2$ we have

$$\mathbb{P}\left(\tau_{\text{inv}}(t) \geq 2ct\right) \leq 6\exp(-t/400\sigma^2).$$

Using union bound, we derive for $t \geq 230\Lambda^2$,

$$\mathbb{P}\left(\tau_{\text{inv}}(t) \geq 2ct \text{ or } \inf_i L^i_{\tau_{\text{inv}}(t)} \leq t/2\right) \leq 12\exp(-t/450\sigma^2).$$

But consider when $\tau_{\text{inv}}(t) \leq 2ct$ and $\inf_i L^i_{\tau_{\text{inv}}(t)} \geq t/2$. This means that by $2ct$, we should have covered state i at least $c_i L_{\tau_{\text{inv}}(t)} \geq c_i t/2$ times (in the continuous sense). We let $t' = 2ct \geq 460c\Lambda^2$, then by t', we should have covered each state at least $\pi_* t'/4$ times. Take $t' \geq 4k/\pi_*$, then we should have covered each state k times by t', which means $\tau_{\text{cov}}^{(k)} \leq t'$. Thus we have for $t' \geq \max\{4k/\pi_*, 460c\Lambda^2\}$,

$$\mathbb{P}(\tau_{\text{cov}}^{(k)} \geq t') \leq 12\exp(-t'/900c\sigma^2) \leq 12\exp(-t'/6000c\Lambda^2).$$

The last step is due to $\sigma^2 \leq 2\pi\Lambda^2$ Ding et al. (2011) (Equation 22). To this end, we have

$$\begin{aligned} t_{\text{cov}}^{(k)} = \mathbb{E}\,\tau_{\text{cov}}^{(k)} &\leq \int_0^{\max\{4k/\pi_*, 460c\Lambda^2\}} \mathbf{1} dt' + \int_{\max\{4k/\pi_*, 460c\Lambda^2\}}^{\infty} \mathbb{P}(\tau_{\text{cov}}^{(k)} \geq t') dt' \\ &\leq \max\{4k/\pi_*, 460c\Lambda^2\} + 80000c\Lambda^2 \\ &\leq 4k/\pi_* + 90000c\Lambda^2 \end{aligned} \tag{8}$$

Thanks to Theorem 3, we have $t_{\text{cov}} = \Theta(c\Lambda^2)$, thus $t_{\text{cov}}^{(k)} = O(k/\pi_* + t_{\text{cov}})$ for continuous chains. ■

To adapt this result for discrete-time Markov chains, we need to use concentration results for sums of i.i.d. exponential random variables.

Lemma 17 (Concentration of exponential RVs.) *Let $\tau_1, \tau_2, ..., \tau_m$ be i.i.d. exponential variables from $Exp(1)$, then the sum of these random variables $S_m = \sum_{i=1}^m \tau_i$ has the following tail concentration bound for $\epsilon \in (0, 1)$.*

$$\mathbb{P}(S_m \geq (1+\epsilon)m) \leq \exp(-m\epsilon^2/4).$$

Proof Note that for any $t > 0$, we have $\mathbb{E}(e^{tS_m}) = (1-t)^{-m}$, thus by Markov's inequality for $\epsilon > 0$, we have $\mathbb{P}(S_m \geq (1+\epsilon)m) \leq \exp(-(1+\epsilon)mt - m\ln(1-t))$. Taking $t = \epsilon/(1+\epsilon)$, we have for $\epsilon \in (0,1)$,

$$\mathbb{P}(S_m \geq (1+\epsilon)m) \leq \exp(\ln(1+\epsilon)m - \epsilon m) \leq \exp(-m\epsilon^2/4).$$

■

We will also use the following lemma proved by Ding et al. (2011) (Lemma 2.4) using the method of majorizing measures.

Lemma 18 (Tail bound summing lemma.) *For random walk over a reversible chain, there exist constant $a, b, u_0 > 0$, such that for any $u \geq u_0$, we have 0-*

$$\sum_{i\in[n]} e^{-uc_i\Lambda^2} \leq ae^{-bu}.$$

Now we are able to translate the result for continuous-time chains to that for discrete-time chains.

Theorem 19 (k-cover time of discrete-time reversible chains) *For discrete-time random walk on reversible chains, we have $t_{\text{cov}}^{(k)} = O(k/\pi_* + t_{\text{cov}})$.*

Proof Fixing any state $i_0 \in [n]$, let $\tau_{inv}(t)$ be the inverse local time for state i_0 of the continuous-time Markov chain. By the proof for Theorem 16, we have for $t \geq 230\Lambda^2$,

$$\mathbb{P}\left(\tau_{\text{inv}}(t) \geq 2ct \text{ or } \inf_i L^i_{\tau_{\text{inv}}(t)} \leq t/2\right) \leq 12\exp(-t/6000\Lambda^2).$$

This means that w.h.p., we have $\tau_{\text{inv}}(t) \leq 2ct$ and we have spent continuous time $c_i t/2$ at state i. However, the probability of taking significantly less jumps in the corresponding discrete Markov chain and get $c_i t/2$ continuous time is very low. Concretely, we have

$$\mathbb{P}(L^i_{\tau_{\text{inv}}(t)} \geq t/2 \mid N_i(\tau_{\text{inv}}(t)) \leq c_i t/4) \leq \exp(-c_i t/16).$$

Denote $E = \{\inf_i L^i_{\tau_{\text{inv}}(t)} \leq \delta t\}$ and $E' \triangleq \{\inf_i \frac{1}{c_i} N_i(\tau_{\text{inv}}(t)) \leq t/4\}$, then we have

$$\mathbb{P}(E \cup E') = \mathbb{P}(E) + \mathbb{P}(E' \cap E^c).$$

But note that for event $E' \cap E^c$, we have $\inf_i \frac{1}{c_i} N_i(\tau_{\text{inv}}(t)) \leq t/4$ and $\inf_i L^i_{\tau_{\text{inv}}(t)} \geq t/2$, so there exists some $i_1 \in [n]$ that satisfies $N_{i_1}(\tau_{\text{inv}}(t)) \leq c_{i_0} t/4$, but $L^{i_1}_{\tau_{\text{inv}}(t)} \geq t/2$. We then use union bound to deduce that

$$\begin{aligned}
\mathbb{P}(E' \cap E^c) &\leq \mathbb{P}(\exists i \in [n], N_i(\tau_{\text{inv}}(t)) \leq c_i t/4 \text{ and } L^i_{\tau_{\text{inv}}(t)} \geq t/2) \\
&\leq \sum_{i=1}^n \mathbb{P}(N_i(\tau_{\text{inv}}(t)) \leq c_i t/4 \text{ and } L^i_{\tau_{\text{inv}}(t)} \geq t/2) \\
&\leq \sum_{i=1}^n \mathbb{P}(L^i_{\tau_{\text{inv}}(t)} \geq t/2 \mid N_i(\tau_{\text{inv}}(t)) \leq c_i t/4) \\
&\leq \sum_{i=1}^n \exp(-c_i t/16).
\end{aligned} \tag{9}$$

Now we can use the tail bound summing lemma (Theorem 18) to deduce that for $t \geq t_0\Lambda^2$ for some constant $t_0 > 0$,

$$\mathbb{P}(E' \cap E^c) \leq a\exp(-bt/\Lambda^2).$$

Here $a, b > 0$ are also constants. Similarly, we define bad events $\Sigma \triangleq \{\tau_{\text{inv}}(t) \geq 2ct\}$ and $\Sigma' \triangleq \{N(\tau_{\text{inv}}(t)) \geq 4ct\}$. Here $N(\tau_{\text{inv}}(t)) = \sum_{i=1}^n N_i(\tau_{\text{inv}}(t))$ is the total number of jumps made before stopping. These bad events happen with probability

$$\mathbb{P}(\Sigma' \cup \Sigma) = \mathbb{P}(\Sigma) + \mathbb{P}(\Sigma'|\Sigma^c)\,\mathbb{P}(\Sigma^c) \leq \mathbb{P}(\Sigma) + \mathbb{P}(\Sigma'|\Sigma^c).$$

Conditioned on the random variable $\tau_{\text{inv}}(t)$, the distribution of $N(\tau_{\text{inv}}(t))$ is Poisson with mean $\tau_{\text{inv}}(t)$ Zhai et al. (2018) (Remark 1.2). By tail bounds for Poisson distribution, we have $\forall x > 0$,

$$\mathbb{P}(N(\tau_{\text{inv}}(t)) \geq \tau_{\text{inv}}(t) + x \mid \tau_{inv}(t) \leq 2ct) \leq \exp\Big(-\frac{x^2}{2(\tau_{\text{inv}}(t)+x)}\Big) \leq \exp\Big(-\frac{x^2}{2(2ct+x)}\Big).$$

When Σ^c is true, Σ' implies $N(\tau_{\text{inv}}(t)) - \tau_{\text{inv}}(t) \geq 4ct - 2ct = 2ct$, which means $N(\tau_{\text{inv}}(t)) \geq \tau_{\text{inv}}(t) + 2ct$. Hence using the tail bound above, we have

$$\mathbb{P}(\Sigma'|\Sigma^c) \leq \mathbb{P}(N(\tau_{\text{inv}}(t)) \geq \tau_{\text{inv}}(t) + 2ct \mid \tau_{inv}(t) \leq 2ct) \leq e^{-ct/2}.$$

Note that $\Omega(n) = t_{\text{cov}} \asymp c\Lambda^2$, therefore we have $\Lambda^{-2} = O(c/n) = o(c)$. By union bound, the bad events $E \cup E' \cup \Sigma \cup \Sigma'$ occurs with probability less than $a'\exp(-b't/\Lambda^2)$ for $t \geq t_0'\Lambda^2$ and some constant $a', b', t_0' > 0$.

However, when no bad event happens, denote $t' := 4ct$ and let $t' \geq 16k/\pi_*$, we have for any $i \in [n]$,

$$N_i(\tau_{\text{inv}}(t)) \geq \frac{c_i t}{4} = \frac{c_i t'}{16c} \geq \frac{c_i k}{c\pi_*} = \frac{\pi_i k}{\pi_*} \geq k.$$

We also note that $N(\tau_{\text{inv}}(t)) \leq t'$, and therefore $\tau_{\text{cov}}^{(k)} \leq t'$. In conclusion, we have shown that for $t' \geq \max\{16k/\pi_*, 4t_0'c\Lambda^2\}$,

$$\mathbb{P}(\tau_{\text{cov}}^{(k)} \geq t') \leq a'\exp(-b't'/4c\Lambda^2).$$

These directly yields that $t_{\text{cov}}^{(k)} = O(k/\pi_* + c\Lambda^2) = O(k/\pi_* + t_{\text{cov}})$. ■

Specially, this gives the tight asymptotic k-cover time for graph random walk. Some interesting instances are as follows.

Example 1 (k-cover time for graph random walks and independent stochastic processes) *We have the following consequences of the theorem above.*

1. *For k-coupon collector, the k-cover time is $t_{\text{cov}}^{(k)} = \Theta(kn + n\ln n)$. The same is true for all regular expanders including the hypercube.*
2. *For full binary tree, the k-cover time is $t_{\text{cov}}^{(k)} = \Theta(kn + n(\ln n)^2)$.*
3. *For cycle and path, the k-cover time is $t_{\text{cov}}^{(k)} = \Theta(kn + n^2)$.*

4. *For non-uniform coupon collector with* $p = (p_1, ..., p_n)$ *and* $p_* = \min_{i\in[n]} p_i$*, the* k*-cover time is* $t_{\text{cov}}^{(k)} = \Theta(k/p_* + t_{\text{cov}})$*.*

However, this only shows that our lower bound is tight for reversible chains. For the general irreducible chains, the isomorphism theorem does not hold and the above arguments cannot be applied.

4.3. Learning and Testing Reversible Chains

In this section, we will see how the k-cover time bound together with previous results on testing/learning discrete distributions together yields sample complexity bounds on learning/testing Markov chains. Specially, we consider Markov chains drawn from the family of chains with cover time upper bounded by t_{cov} and minimum stationary probability lower bounded by π_*, and we denote this family as $\mathcal{M}_{rev}(t_{\text{cov}}, \pi_*)$.

We have the following theorem on testing and learning Markov chains due to theorems and lemmas proved thus far.

Theorem 20 (Sample complexity bounds for learning/testing reversible chains) *For a* n*-state reversible Markov chains from* $\mathcal{M}_{rev}(t_{\text{cov}}, \pi_*)$*, we have the following sample complexity bounds.*

1. *We can* (ϵ, δ)*-learn the chain using* $O_\delta(t_{\text{cov}} + \frac{n \ln n}{\pi_* \epsilon^2})$ *samples;*
2. *We can* (ϵ, δ)*-uniform-test the chain using* $O_\delta(n \ln n + \frac{\sqrt{n} \ln n}{\pi_* \epsilon^2})$ *samples;*
3. *We can* (ϵ, δ)*-identity-test the chain using* $O_\delta(t_{\text{cov}} + \frac{\sqrt{n} \ln n}{\pi_* \epsilon^2})$ *samples;*
4. *We can* (ϵ, δ)*-closeness-test the chains using* $O_\delta(t_{\text{cov}} + \frac{\ln n}{\pi_*}(\frac{n^{2/3}}{\epsilon^{4/3}} + \frac{\sqrt{n}}{\epsilon^2}))$ *samples.*
5. *We can* $(\epsilon_1, \epsilon_2, \delta)$*-tolerant-uniform/identity/closeness-test the chain using* $O_\delta(t_{\text{cov}} + \frac{n}{\pi_*(\epsilon_2 - \epsilon_1)^2})$ *samples.*
6. *We can* $(\epsilon/2\sqrt{n}, \epsilon, \delta)$*-tolerant-uniform-test the chain using* $O_\delta(t_{\text{cov}} + \frac{\sqrt{n} \ln n}{\pi_* \epsilon^4})$ *samples.*
7. *We can* $(\epsilon^3/300\sqrt{n} \ln n, \epsilon, \delta)$*-tolerant-identity-test the chain using* $O_\delta(t_{\text{cov}} + \frac{\sqrt{n} \ln n}{\pi_* \epsilon^6})$ *samples.*

Proof This is a direct application of Theorem 1, Theorem 16 and Theorem 6, Theorem 7, Theorem 8. For example, the sample complexity of learning Markov chain M is $O_\delta(t_{\text{cov}}^{k(n,\epsilon,\delta/2n)}) = O_\delta(t_{\text{cov}} + k(n, \epsilon, \delta/2n)/\pi_*) = O_\delta(t_{\text{cov}} + n \ln n/\pi_* \epsilon^2)$. ■

5. The k-cover Time of Irreducible Chains

For general irreducible chains, the connections with resistance network and Gaussian free field no longer hold. However, we can still use the bounds of k-return time for irreducible chains to bound the k-cover time up to logarithmic factors. We conjecture that the lower bound is tight, i.e., we have $t_{\text{cov}}^{(k)} = \Theta(t_{\text{cov}} + k/\pi_*)$ for all irreducible chains, but we believe advanced tools will be needed to prove this conjecture.

5.1. Upper Bound for Irreducible Chains

The proof of the tight upper bound on k-cover time for reversible chains uses the connections between the cover time and effective resistance/Gaussian free field, none of which have a nice counterpart for general irreducible chains. However, one can still prove upper bounds on the k-cover time that's $O(\ln n)$-factor looser, by bounding the k-hitting times (i.e. the first time when a particular state i_0 is visited k times). Similar idea is used to derive upper bounds on the blanket time in Winkler and Zuckerman (1996).

Lemma 21 (Concentration of the k-hitting time) *For random walk on irreducible chains, the k-hitting time of state $i \in V$ satisfies*

$$\mathbb{P}(\tau_{\text{hit}}^{(k)}(i) \geq t) \leq e \cdot \exp(-t/e(t_{\text{hit}} + k/\pi_i)),$$

for any $t \geq 0$.

Proof Note for irreducible chains, we still have $t_{\text{ret}}(i) = 1/\pi_i$, and hence $t_{\text{hit}}^{(k)}(i) = t_{\text{hit}}(i) + (k-1)/\pi_i \leq t_{\text{hit}} + (k-1)/\pi_i$ by the Markov property. Hence $\mathbb{P}(\tau_{\text{hit}}^{(k)}(i) \geq e(t_{\text{hit}} + (k-1)/\pi_i)) \leq 1/e$. By similar argument used in the exponential decay lemma, we have $\mathbb{P}(\tau_{\text{hit}}^{(k)}(i) \geq em(t_{\text{hit}} + (k-1)/\pi_i)) \leq 1/e^m$, and therefore

$$\mathbb{P}(\tau_{\text{hit}}^{(k)}(i) \geq t) \leq e \cdot \exp(-t/e(t_{\text{hit}} + (k-1)/\pi_i)) \leq e \cdot \exp(-t/e(t_{\text{hit}} + k/\pi_i)).$$

This proves the lemma. ∎

This yields an upper bound on k-cover time for irreducible chains as the following.

Theorem 22 (k-cover time from k-hitting time) *For random walk on irreducible chains, we have $t_{\text{cov}}^{(k)} = \tilde{O}(t_{\text{cov}} + k/\pi_*)$.*

Proof We can think of the k-cover time $\tau_{\text{cov}}^{(k)}$ upper bounded by the maximum of k-hitting times of different states. That is, we have $\tau_{\text{cov}}^{(k)} \leq \max_{i \in V} \tau_{\text{hit}}^{(k)}(i)$. Then by the concentration of k-hitting time, we have $\mathbb{P}(\tau_{\text{hit}}^{(k)}(i) \geq t) \leq e \cdot \exp(-t/e(t_{\text{hit}} + k/\pi_i))$. By a union bound, we have

$$\mathbb{P}(\max_{i \in V} \tau_{\text{hit}}^{(k)}(i) \geq t) \leq \sum_{i \in [n]} e \cdot \exp(-t/e(t_{\text{hit}} + k/\pi_i)) \leq en \cdot \exp(-t/e(t_{\text{hit}} + k/\pi_*)).$$

Now we can transform this high probability bound into expectation form as

$$\begin{aligned} \mathbb{E}[\tau_{\text{cov}}^{(k)}] \leq & \int_{t=0}^{e \ln n (t_{\text{hit}} + k/\pi_*)} 1 \cdot dt + \int_{t = e \ln n (t_{\text{hit}} + k/\pi_*)}^{\infty} en \cdot \exp(-t/e(t_{\text{hit}} + k/\pi_*)) \cdot dt \\ = & e \ln n (t_{\text{hit}} + k/\pi_*) + e^2 (t_{\text{hit}} + k/\pi_*), \end{aligned} \tag{10}$$

This gives us that $t_{\text{cov}}^{(k)} = O(t_{\text{cov}} \ln n + k \ln n/\pi_*) = \tilde{O}(t_{\text{cov}} + k/\pi_*)$. ∎

5.2. Upper Bound for Ergodic Chains

In Wolfer and Kontorovich (2019b), the family of ergodic chains with pseudo-spectral gap lower bounded by γ_{ps} and minimum stationary probability lower bounded by π_* as $\mathcal{M}_{erg}(\gamma_{\text{ps}}, \pi_*)$ is considered.

We remark that this is a sub-family of irreducible chains that have finite-time mixing properties. It excludes all periodic random walks, including simple random walk on a two-node single-edge graph. Thus, our arguments via the k-cover time in fact broaden the family of chains that previous results can be applied to. Specially, we use Paulin's result Paulin (2015) to bound the k-cover time w.r.t. the pseudo-spectral gap γ_{ps}. This naturally recovers the results in Wolfer and Kontorovich (2019b, 2020a).

For an irreducible Markov chain X_1^m over $[n]$, given a function $f : [n] \to \mathbf{R}$ satisfying $f \in L_2(\pi)$, i.e., $\mathbb{E}_{i\sim\pi} f^2(i) = \sum_{i=1}^n \pi_i f^2(i) < \infty$, then it will have finite *stationary expectation* as $E_f \triangleq \mathbb{E}_\pi f = \mathbb{E}_{i\sim\pi} f(i) < \infty$ and finite *stationary variance* as $V_f \triangleq \mathbf{Var}_\pi(f) = \mathbb{E}_\pi f^2 - (\mathbb{E}_\pi f)^2 < \infty$. Then we have the following concentration inequality over Markov chains due to Paulin (2015).

Lemma 23 (Bernstein inequality for Markov chains) *For an irreducible Markov chain X_1^m over $[n]$ and given $f \in L_2(\pi)$, if we have $|f(i) - \mathbb{E}_\pi(f)| \le C, \forall i \in [n]$, and let $S = \sum_{i=1}^m f(X_i)$, then for any starting distribution,*

$$\mathbb{P}\left(\left|S - \mathbb{E}_\pi(S)\right| \ge t\right) \le \sqrt{\frac{2}{\pi_*}} \exp\left(\frac{-t^2\gamma_{\text{ps}}}{-16(m + 1/\gamma_{\text{ps}})V_f + 40tC}\right).$$

Similar inequality is true for reversible Markov chains, with γ_* instead of γ_{ps} in the right hand side and is slightly tighter.

Lemma 24 (High probability bound on k-cover time of ergodic chain) *For an ergodic Markov chain with minimum stationary probability π_* and pseudo-spectral gap γ_{ps}, when $m \ge \max\{\frac{300}{\pi_*\gamma_{\text{ps}}} \ln\left(\frac{n}{\delta}\sqrt{\frac{2}{\pi_*}}\right), \frac{2k}{\pi_*}\}$, we have $\mathbb{P}(\{\tau_{\text{cov}} \le m\}) \ge 1 - \delta$.*

Proof Denote the event $E_i \triangleq \{N_i(m) \in [0.5m\pi_i, 1.5m\pi_i]\}$, and event $E \triangleq \cup_{i\in[n]} E_i$. Then due to Paulin's result, we have

$$\mathbb{P}(E_i) = \mathbb{P}(|N_i(m) - m\pi_i| \ge 0.5m\pi_i) \le \sqrt{\frac{2}{\pi_*}} \exp\Big(-\frac{m^2\pi_i\gamma_{\text{ps}}}{64(m + 1/\gamma_{\text{ps}}) + 80m}\Big)$$

Here we used $f(j) = \delta_i^j, \forall j \in [n]$, then $|f(j) - \mathbb{E}_\pi(f)| \le 1$, and $V_f = \pi_i(1 - \pi_i) \le \pi_i, S = N_i(m) = \sum_{i=1}^m f(X_i)$. Then by a union bound, we have

$$\begin{aligned} \mathbb{P}(E) \le \sum_{x\in[n]} \mathbb{P}(E_x) \le \sum_{i\in[n]} \sqrt{\frac{2}{\pi_*}} \exp\Big(-\frac{m^2\pi_i\gamma_{\text{ps}}}{64(m + 1/\gamma_{\text{ps}}) + 80m}\Big) \\ \le n\sqrt{\frac{2}{\pi_*}} \exp\Big(-\frac{m^2\pi_*\gamma_{\text{ps}}}{150(m + 1/\gamma_{\text{ps}})}\Big). \end{aligned} \tag{11}$$

Thus if we denote $\alpha_\delta \triangleq \frac{150}{\pi_*} \ln\left(\frac{n}{\delta}\sqrt{\frac{2}{\pi_*}}\right)$, then we have for $m \ge \frac{\alpha_\delta}{2\gamma_{\text{ps}}} + \frac{1}{2}\sqrt{\left(\frac{\alpha_\delta}{\gamma_{\text{ps}}}\right)^2 + 4\frac{\alpha_\delta}{\gamma_{\text{ps}}^2}}$, $\mathbb{P}(E) \le \delta$. Note $\alpha_\delta > 1$, and use $\sqrt{x + y} \le \sqrt{x} + \sqrt{y}$, we have that $m \ge 2\frac{\alpha_\delta}{\gamma_{\text{ps}}}$ suffices to make

$\mathbb{P}(E) \leq \delta$. Note we have $\mathbb{P}(E^c) \leq P(\{\forall i \in [n], N_i(m) \geq 0.5m\pi_x\}) \leq \mathbb{P}(\{\forall i \in [n], N_i(m) \geq 0.5m\pi_*\}) \leq \mathbb{P}(\{\tau_{\text{cov}}^{(0.5m\pi_*)} \leq m\})$. Let $m \geq \frac{2k}{\pi_*}$, then $\mathbb{P}(\{\tau_{\text{cov}}^{(k)} \leq m\}) \geq \mathbb{P}(\{\tau_{\text{cov}}^{(0.5m\pi_*)} \leq m\})$. Thus, for $m \geq \max\{\frac{2k}{\pi_*}, 2\frac{\alpha_\delta}{\gamma_{\text{ps}}}\}$, we have $\mathbb{P}(\{\tau_{\text{cov}}^{(k)} \leq m\}) \geq 1-\delta$. ■

This then gives the following bound on the expected k-cover time.

Theorem 25 (Upper bound on expected k-cover time for irreducible chain) *For an ergodic Markov chain with π_* and pseudo-spectral gap γ_{ps}, we have $t_{\text{cov}}^{(k)} \leq \max\{\frac{4k}{\pi_*}, \frac{600}{\pi_*\gamma_{\text{ps}}} \ln\left(\frac{150\sqrt{2}n}{\sqrt{\pi_*}}\right)\}$, and hence $t_{\text{cov}}^{(k)} = O(\frac{k}{\pi_*} + \frac{1}{\pi_*\gamma_{\text{ps}}} \ln \frac{n}{\pi_*})$.*

Proof Since for $m \geq \max\{\frac{300}{\pi_*\gamma_{\text{ps}}} \ln\left(\frac{n}{\delta}\sqrt{\frac{2}{\pi_*}}\right), \frac{2k}{\pi_*}\}$, we have $\mathbb{P}(\{\tau_{\text{cov}} \leq m\}) \geq 1-\delta$. Thus let $t = \frac{300}{\pi_*\gamma_{\text{ps}}} \ln\left(\frac{n}{\delta}\sqrt{\frac{2}{\pi_*}}\right)$, we have $\mathbb{P}(\tau_{\text{cov}}^{(k)} \leq \max\{t, 2k/\pi_*\}) \geq 1 - n\sqrt{\frac{2}{\pi_*}} \exp\left(-\frac{\pi_*\gamma_{\text{ps}}t}{300}\right)$. Thus, the expected k-cover time

$$\begin{aligned} \mathbb{E}(\tau_{\text{cov}}^{(k)}) &= \int_0^{\frac{2k}{\pi_*}} \mathbb{P}(\tau_{\text{cov}}^{(k)} \geq t)dt + \int_{\frac{2k}{\pi_*}}^{\infty} \mathbb{P}(\tau_{\text{cov}}^{(k)} \geq t)dt \\ &\leq \frac{2k}{\pi_*} + \int_{\frac{2k}{\pi_*}}^{\infty} n\sqrt{\frac{2}{\pi_*}} \exp\left(-\frac{\pi_*\gamma_{\text{ps}}t}{300}\right)dt. \end{aligned} \tag{12}$$

This finally gives $t_{\text{cov}}^{(k)} \leq \frac{2k}{\pi_*} + \frac{300\sqrt{2}n}{\gamma_{\text{ps}}\pi_*^{3/2}} \exp\left(-\frac{\gamma_{\text{ps}}k}{150}\right)$. When $k \geq k^* \triangleq \frac{150}{\gamma_{\text{ps}}} \ln\left(\frac{150\sqrt{2}n}{\sqrt{\pi_*}}\right)$, we have $\frac{150\sqrt{2}n}{\sqrt{\pi_*}} \exp\left(-\frac{\gamma_{\text{ps}}k}{150}\right) \leq 1$, and therefore, $\frac{300\sqrt{2}n}{\gamma_{\text{ps}}\pi_*^{3/2}} \exp\left(-\frac{\gamma_{\text{ps}}k}{150}\right) \leq \frac{2}{\gamma_{\text{ps}}\pi_*}$. Note $k \geq \frac{1}{\gamma_{\text{ps}}}$, thus $\frac{2k}{\pi_*} \geq \frac{300\sqrt{2}n}{\gamma_{\text{ps}}\pi_*^{3/2}} \exp\left(-\frac{\gamma_{\text{ps}}k}{150}\right)$ and hence $t_{\text{cov}}^{(k)} \leq \frac{4k}{\pi_*}$. But when $k \leq k^*$, we have $t_{\text{cov}}^{(k)} \leq t_{\text{cov}}^{(k^*)} \leq \frac{4k^*}{\pi_*}$. This proves the theorem. ■

Remark 26 (Concentration inequality for general irreducible chains) *In Moulos (2020), the following concentration inequality is proved. Here $f : [n] \to (a,b)$ is any bounded function on the state space and $\mathbf{q}$ is the initial distribution. For any irreducible Markov chain,*

$$\mathbb{P}\left(\left|S - \mathop{\mathbb{E}}_{\mathbf{q}}(S)\right| \geq t\right) \leq \sqrt{\frac{2}{\pi_*}} \exp\left(\frac{-t^2}{2m(b-a)^2 t_{\text{hit}}^2}\right).$$

However, since the right hand side incurs a quadratic dependence on $t_{\text{hit}}^2 = \tilde{\Theta}(t_{\text{cov}}^2)$, it would yield a much worse bound on the k-cover time than that in Theorem 22.

5.3. Learning and Testing Irreducible (or Ergodic) Chains

Again, we will see how the k-cover time bound implies sample complexity bounds on learning/testing Markov chains. We consider the family of irreducible chains with cover time upper bounded by t_{cov} and minimum stationary probability lower bounded by π_*, which we denote as $\mathcal{M}_{irr}(t_{\text{cov}}, \pi_*)$. We also consider $\mathcal{M}_{erg}(\gamma_{\text{ps}}, \pi_*)$, the family of ergodic chains with pseudo-spectral gap lower bounded by γ_{ps} and minimum stationary probability

The following theorem on testing and learning Markov chains is a natural corollary of the theorems proved.

Theorem 27 (Sample complexity bounds for learning/testing irreducible chains) *For a n-state irreducible Markov chains from $\mathcal{M}_{irr}(t_{\text{cov}}, \pi_*)$ (or $\mathcal{M}_{erg}(\gamma_{\text{ps}}, \pi_*)$), we have the following sample complexity bounds.*

1. *We can (ϵ, δ)-learn the chain using $\tilde{O}(t_{\text{cov}} + \frac{n}{\pi_* \epsilon^2})$ (or $\tilde{O}(\frac{1}{\pi_* \gamma_{\text{ps}}} + \frac{n}{\pi_* \epsilon^2})$) samples;*
2. *We can (ϵ, δ)-identity-test the chain using $\tilde{O}(t_{\text{cov}} + \frac{\sqrt{n}}{\pi_* \epsilon^2})$ (or $\tilde{O}(\frac{1}{\pi_* \gamma_{\text{ps}}} + \frac{\sqrt{n}}{\pi_* \epsilon^2})$) samples;*
3. *We can (ϵ, δ)-closeness-test the chains using $\tilde{O}(t_{\text{cov}} + \frac{1}{\pi_*}(\frac{n^{2/3}}{\epsilon^{4/3}} + \frac{\sqrt{n}}{\epsilon^2}))$ (or $\tilde{O}(\frac{1}{\pi_* \gamma_{\text{ps}}} + \frac{1}{\pi_*}(\frac{n^{2/3}}{\epsilon^{4/3}} + \frac{\sqrt{n}}{\epsilon^2}))$) samples.*
4. *The other results are analogous to Theorem 20, within a logarithmic factor.*

Proof This is a direct application of Theorem 1, Theorem 22 and Theorem 6, Theorem 7, Theorem 8. Results about Markov chains from $\mathcal{M}_{ierg}(\gamma_{\text{ps}}, \pi_*)$ uses Theorem 25 instead of Theorem 22. ■

6. Conclusion and Open Problems

In this paper, we considered the problem of testing and learning Markov chains from a single trajectory. We show that the sample complexity of a number of learning and testing problems over Markov chains is strongly related to the k-cover time of the unknown chain. We then proved that $t_{\text{cov}}^{(k)} = \Theta(t_{\text{cov}} + k/\pi_*)$ for reversible Markov chains and $t_{\text{cov}}^{(k)} = \tilde{\Theta}(t_{\text{cov}} + k/\pi_*)$ for general irreducible Markov chains. These results on k-cover time give sample complexity bounds for a broad family of learning and testing problems over Markov chains, and apply to a broader family of chains than in the previous works.

We leave the tight characterization of k-cover time for irreducible chains as an open problem, but we conjecture the lower bound to be tight. It would also be nice if one can prove corresponding lower bounds on sample complexity using the idea of k-cover times.

Moreover, it has been considered by Newman that the second set of coupon in the coupon collector's problem costs $\Theta(n \ln \ln n)$, even though the first set of coupon costs $\Theta(n \ln n)$ in expectation. It's dubbed the "double dixie cup problem" in Newman (1960). We find it interesting to ask similar questions in the setting of Markov chains: what's the cost of a second cover in n-cycle, n-path or torus? By our theorem for k-cover time, for k large enough, it seems that each marginal cover costs $\Theta(1/\pi_*)$, but it gives no clue about the cost of the second cover.

References

David Aldous and James Fill. Reversible markov chains and random walks on graphs, 1995.

Martin Anthony and Peter L Bartlett. *Neural Network Learning: Theoretical Foundations*. cambridge university press, 2009.

Tugkan Batu, Eldar Fischer, Lance Fortnow, Ravi Kumar, Ronitt Rubinfeld, and Patrick White. Testing random cariables for independence and identity. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 442–451. IEEE, 2001.

Clément L Canonne. A survey on distribution testing. 2017.

Siu-On Chan, Ilias Diakonikolas, Paul Valiant, and Gregory Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1193–1203. SIAM, 2014.

Yeshwanth Cherapanamjeri and Peter L Bartlett. Testing symmetric markov chains without hitting. In *Conference on Learning Theory*, pages 758–785, 2019.

Constantinos Daskalakis, Nishanth Dikkala, and Nick Gravin. Testing symmetric markov chains from a single trajectory. *arXiv preprint arXiv:1704.06850*, 2017.

Jian Ding. Asymptotics of cover times via gaussian free fields: Bounded-degree graphs and general trees. *The Annals of Probability*, 42(2):464–496, 2014.

Jian Ding, James R Lee, and Yuval Peres. Cover times, blanket times, and majorizing measures. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 61–70, 2011.

Nathalie Eisenbaum, Haya Kaspi, Michael B. Marcus, Jay Rosen, and Zhan Shi. A Ray-Knight theorem for symmetric Markov processes. *Annals of Probability*, 28(4):1781–1796, 2000.

Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 68–75. Springer, 2011.

Yi Hao, Alon Orlitsky, and Venkatadheeraj Pichapati. On learning markov chains. In *Advances in Neural Information Processing Systems*, pages 648–657, 2018.

Daniel J Hsu, Aryeh Kontorovich, and Csaba Szepesvári. Mixing time estimation in reversible markov chains from a single sample path. In *Advances in neural information processing systems*, pages 1459–1467, 2015.

Aryeh Kontorovich and Gergely Neu, editors. *Algorithmic Learning Theory, ALT 2020, 8-11 February 2020, San Diego, CA, USA*, volume 117 of *Proceedings of Machine Learning Research*, 2020. PMLR. URL http://proceedings.mlr.press/v117/.

David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

M. B. Marcus and J. Rosen. *Markov processes, Gaussian processes, and local times*. Cambridge, 2006.

Vrettos Moulos. A hoeffding inequality for finite state markov chains and its applications to markovian bandits. In *IEEE International Symposium on Information Theory*, 2020.

Donald J. Newman. The double dixie cup problem. *The American Mathematical Monthly*, 1960.

L. Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Trans. Inf. Theor.*, 54(10):4750–4755, 2008.

Daniel Paulin. Concentration inequalities for markov chains by marton couplings and spectral methods. *Electronic Journal of Probability*, 20, 2015.

Gregory Valiant and Paul Valiant. The power of linear estimators. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 403–412. IEEE, 2011.

Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.

Peter Winkler and David Zuckerman. Multiple cover time. *Random Structures & Algorithms*, 9(4): 403–411, 1996.

Geoffrey Wolfer. Mixing time estimation in ergodic markov chains from a single trajectory with contraction methods. *Algorithmic Learning Theory*, 2020.

Geoffrey Wolfer and Aryeh Kontorovich. Estimating the mixing time of ergodic markov chains. *Conference on Learning Theory*, 2019a.

Geoffrey Wolfer and Aryeh Kontorovich. Minimax learning of ergodic markov chains. In *Algorithmic Learning Theory*, pages 904–930, 2019b.

Geoffrey Wolfer and Aryeh Kontorovich. Minimax testing of identity to a reference rrgodic markov chain. *International Conference on Artificial Intelligence and Statistics*, 2020a.

Geoffrey Wolfer and Aryeh Kontorovich. Statistical estimation of ergodic markov chain kernel over discrete state space. *ArXiv:1809.05014v4*, 2020b.

Alex Zhai et al. Exponential concentration of cover times. *Electronic Journal of Probability*, 23, 2018.