
Mixing Time Estimation in Reversible Markov Chains from a Single Sample Path

Daniel Hsu
Columbia University
djhsu@cs.columbia.edu

Aryeh Kontorovich
Ben-Gurion University
karyeh@cs.bgu.ac.il

Csaba Szepesvári
University of Alberta
szepesva@cs.ualberta.ca

Abstract

This article provides the first procedure for computing a fully data-dependent interval that traps the mixing time t_{mix} of a finite reversible ergodic Markov chain at a prescribed confidence level. The interval is computed from a single finite-length sample path from the Markov chain, and does not require the knowledge of any parameters of the chain. This stands in contrast to previous approaches, which either only provide point estimates, or require a reset mechanism, or additional prior knowledge. The interval is constructed around the relaxation time t_{relax} , which is strongly related to the mixing time, and the width of the interval converges to zero roughly at a \sqrt{n} rate, where n is the length of the sample path. Upper and lower bounds are given on the number of samples required to achieve constant-factor multiplicative accuracy. The lower bounds indicate that, unless further restrictions are placed on the chain, no procedure can achieve this accuracy level before seeing each state at least $\Omega(t_{\text{relax}})$ times on the average. Finally, future directions of research are identified.

1 Introduction

This work tackles the challenge of constructing fully empirical bounds on the mixing time of Markov chains based on a single sample path. Let $(X_t)_{t=1,2,\dots}$ be an irreducible, aperiodic time-homogeneous Markov chain on a finite state space $[d] := \{1, 2, \dots, d\}$ with transition matrix \mathbf{P} . Under this assumption, the chain converges to its unique stationary distribution $\boldsymbol{\pi} = (\pi_i)_{i=1}^d$ regardless of the initial state distribution \mathbf{q} :

$$\lim_{t \rightarrow \infty} \Pr_{\mathbf{q}}(X_t = i) = \lim_{t \rightarrow \infty} (\mathbf{q}\mathbf{P}^t)_i = \pi_i \quad \text{for each } i \in [d].$$

The *mixing time* t_{mix} of the Markov chain is the number of time steps required for the chain to be within a fixed threshold of its stationary distribution:

$$t_{\text{mix}} := \min \left\{ t \in \mathbb{N} : \sup_{\mathbf{q}} \max_{A \subset [d]} |\Pr_{\mathbf{q}}(X_t \in A) - \boldsymbol{\pi}(A)| \leq 1/4 \right\}. \quad (1)$$

Here, $\boldsymbol{\pi}(A) = \sum_{i \in A} \pi_i$ is the probability assigned to set A by $\boldsymbol{\pi}$, and the supremum is over all possible initial distributions \mathbf{q} . The problem studied in this work is the construction of a non-trivial confidence interval $C_n = C_n(X_1, X_2, \dots, X_n, \delta) \subset [0, \infty]$, based only on the observed sample path (X_1, X_2, \dots, X_n) and $\delta \in (0, 1)$, that succeeds with probability $1 - \delta$ in trapping the value of the mixing time t_{mix} .

This problem is motivated by the numerous scientific applications and machine learning tasks in which the quantity of interest is the mean $\boldsymbol{\pi}(f) = \sum_i \pi_i f(i)$ for some function f of the states of a Markov chain. This is the setting of the celebrated Markov Chain Monte Carlo (MCMC) paradigm [1], but the problem also arises in performance prediction involving time-correlated data, as is common in reinforcement learning [2]. Observable bounds on mixing times are useful in the

design and diagnostics of these methods; they yield effective approaches to assessing the estimation quality, even when *a priori* knowledge of the mixing time or correlation structure is unavailable.

Main results. We develop the first procedure for constructing non-trivial and fully empirical confidence intervals for Markov mixing time. Consider a reversible ergodic Markov chain on d states with absolute spectral gap γ_* and stationary distribution minorized by π_* . As is well-known [3, Theorems 12.3 and 12.4],

$$(t_{\text{relax}} - 1) \ln 2 \leq t_{\text{mix}} \leq t_{\text{relax}} \ln \frac{4}{\pi_*} \quad (2)$$

where $t_{\text{relax}} := 1/\gamma_*$ is the *relaxation time*. Hence, it suffices to estimate γ_* and π_* . Our main results are summarized as follows.

1. In Section 3.1, we show that in some problems $n = \Omega((d \log d)/\gamma_* + 1/\pi_*)$ observations are necessary for any procedure to guarantee constant multiplicative accuracy in estimating γ_* (Theorems 1 and 2). Essentially, in some problems *every* state may need to be visited about $\log(d)/\gamma_*$ times, on average, before an accurate estimate of the mixing time can be provided, regardless of the actual estimation procedure used.
2. In Section 3.2, we give a point-estimator for γ_* , and prove in Theorem 3 that it achieves multiplicative accuracy from *a single sample path* of length $\tilde{O}(1/(\pi_* \gamma_*^3))$.¹ We also provide a point-estimator for π_* that requires a sample path of length $\tilde{O}(1/(\pi_* \gamma_*))$. This establishes the feasibility of *estimating* the mixing time in this setting. However, the valid confidence intervals suggested by Theorem 3 depend on the unknown quantities π_* and γ_* . We also discuss the importance of reversibility, and some possible extensions to non-reversible chains.
3. In Section 4, the construction of valid *fully empirical confidence intervals* for π_* and γ_* are considered. First, the difficulty of the task is explained, i.e., why the standard approach of turning the finite time confidence intervals of Theorem 3 into a fully empirical one fails. Combining several results from perturbation theory in a novel fashion we propose a new procedure and prove that it avoids slow convergence (Theorem 4). We also explain how to combine the empirical confidence intervals from Algorithm 1 with the non-empirical bounds from Theorem 3 to produce valid empirical confidence intervals. We prove in Theorem 5 that the width of these new intervals converge to zero asymptotically at least as fast as those from either Theorem 3 and Theorem 4.

Related work. There is a vast statistical literature on estimation in Markov chains. For instance, it is known that under the assumptions on $(X_t)_t$ from above, the law of large numbers guarantees that the sample mean $\pi_n(f) := \frac{1}{n} \sum_{t=1}^n f(X_t)$ converges almost surely to $\pi(f)$ [4], while the central limit theorem tells us that as $n \rightarrow \infty$, the distribution of the deviation $\sqrt{n}(\pi_n(f) - \pi(f))$ will be normal with mean zero and asymptotic variance $\lim_{n \rightarrow \infty} n \text{Var}(\pi_n(f))$ [5].

Although these asymptotic results help us understand the limiting behavior of the sample mean over a Markov chain, they say little about the finite-time non-asymptotic behavior, which is often needed for the prudent evaluation of a method or even its algorithmic design [6–13]. To address this need, numerous works have developed Chernoff-type bounds on $\Pr(|\pi_n(f) - \pi(f)| > \epsilon)$, thus providing valuable tools for non-asymptotic probabilistic analysis [6, 14–16]. These probability bounds are larger than corresponding bounds for independent and identically distributed (iid) data due to the temporal dependence; intuitively, for the Markov chain to yield a fresh draw $X_{t'}$ that behaves as if it was independent of X_t , one must wait $\Theta(t_{\text{mix}})$ time steps. Note that the bounds generally depend on distribution-specific properties of the Markov chain (e.g., P , t_{mix} , γ_*), which are often unknown *a priori* in practice. Consequently, much effort has been put towards estimating these unknown quantities, especially in the context of MCMC diagnostics, in order to provide data-dependent assessments of estimation accuracy [e.g., 11, 12, 17–19]. However, these approaches generally only provide asymptotic guarantees, and hence fall short of our goal of empirical bounds that are valid with any finite-length sample path.

Learning with dependent data is another main motivation to our work. Many results from statistical learning and empirical process theory have been extended to sufficiently fast mixing, dependent

¹The $\tilde{O}(\cdot)$ notation suppresses logarithmic factors.

data [e.g., 20–26], providing learnability assurances (e.g., generalization error bounds). These results are often given in terms of mixing coefficients, which can be consistently estimated in some cases [27]. However, the convergence rates of the estimates from [27], which are needed to derive confidence bounds, are given in terms of unknown mixing coefficients. When the data comes from a Markov chain, these mixing coefficients can often be bounded in terms of mixing times, and hence our main results provide a way to make them fully empirical, at least in the limited setting we study.

It is possible to eliminate many of the difficulties presented above when allowed more flexible access to the Markov chain. For example, given a sampling oracle that generates independent transitions from any given state (akin to a “reset” device), the mixing time becomes an efficiently testable property in the sense studied in [28, 29]. On the other hand, when one only has a circuit-based description of the transition probabilities of a Markov chain over an exponentially-large state space, there are complexity-theoretic barriers for many MCMC diagnostic problems [30].

2 Preliminaries

2.1 Notations

We denote the set of positive integers by \mathbb{N} , and the set of the first d positive integers $\{1, 2, \dots, d\}$ by $[d]$. The non-negative part of a real number x is $[x]_+ := \max\{0, x\}$, and $\lceil x \rceil_+ := \max\{0, \lceil x \rceil\}$. We use $\ln(\cdot)$ for natural logarithm, and $\log(\cdot)$ for logarithm with an arbitrary constant base. Bold-face symbols are used for vectors and matrices (e.g., \mathbf{v} , \mathbf{M}), and their entries are referenced by subindexing (e.g., v_i , $M_{i,j}$). For a vector \mathbf{v} , $\|\mathbf{v}\|$ denotes its Euclidean norm; for a matrix \mathbf{M} , $\|\mathbf{M}\|$ denotes its spectral norm. We use $\text{Diag}(\mathbf{v})$ to denote the diagonal matrix whose (i, i) -th entry is v_i . The probability simplex is denoted by $\Delta^{d-1} = \{\mathbf{p} \in [0, 1]^d : \sum_{i=1}^d p_i = 1\}$, and we regard vectors in Δ^{d-1} as row vectors.

2.2 Setting

Let $\mathbf{P} \in (\Delta^{d-1})^d \subset [0, 1]^{d \times d}$ be a $d \times d$ row-stochastic matrix for an ergodic (i.e., irreducible and aperiodic) Markov chain. This implies there is a unique stationary distribution $\boldsymbol{\pi} \in \Delta^{d-1}$ with $\pi_i > 0$ for all $i \in [d]$ [3, Corollary 1.17]. We also assume that \mathbf{P} is *reversible* (with respect to $\boldsymbol{\pi}$):

$$\pi_i P_{i,j} = \pi_j P_{j,i}, \quad i, j \in [d]. \quad (3)$$

The minimum stationary probability is denoted by $\pi_\star := \min_{i \in [d]} \pi_i$.

Define the matrices

$$\mathbf{M} := \text{Diag}(\boldsymbol{\pi})\mathbf{P} \quad \text{and} \quad \mathbf{L} := \text{Diag}(\boldsymbol{\pi})^{-1/2}\mathbf{M}\text{Diag}(\boldsymbol{\pi})^{-1/2}.$$

The (i, j) th entry of the matrix $M_{i,j}$ contains the *doublet probabilities* associated with \mathbf{P} : $M_{i,j} = \pi_i P_{i,j}$ is the probability of seeing state i followed by state j when the chain is started from its stationary distribution. The matrix \mathbf{M} is symmetric on account of the reversibility of \mathbf{P} , and hence it follows that \mathbf{L} is also symmetric. (We will strongly exploit the symmetry in our results.) Further, $\mathbf{L} = \text{Diag}(\boldsymbol{\pi})^{1/2}\mathbf{P}\text{Diag}(\boldsymbol{\pi})^{-1/2}$, hence \mathbf{L} and \mathbf{P} are similar and thus their eigenvalue systems are identical. Ergodicity and reversibility imply that the eigenvalues of \mathbf{L} are contained in the interval $(-1, 1]$, and that 1 is an eigenvalue of \mathbf{L} with multiplicity 1 [3, Lemmas 12.1 and 12.2]. Denote and order the eigenvalues of \mathbf{L} as

$$1 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_d > -1.$$

Let $\lambda_\star := \max\{\lambda_2, |\lambda_d|\}$, and define the (absolute) spectral gap to be $\gamma_\star := 1 - \lambda_\star$, which is strictly positive on account of ergodicity.

Let $(X_t)_{t \in \mathbb{N}}$ be a Markov chain whose transition probabilities are governed by \mathbf{P} . For each $t \in \mathbb{N}$, let $\boldsymbol{\pi}^{(t)} \in \Delta^{d-1}$ denote the marginal distribution of X_t , so

$$\boldsymbol{\pi}^{(t+1)} = \boldsymbol{\pi}^{(t)}\mathbf{P}, \quad t \in \mathbb{N}.$$

Note that the initial distribution $\boldsymbol{\pi}^{(1)}$ is arbitrary, and need not be the stationary distribution $\boldsymbol{\pi}$.

The goal is to estimate π_\star and γ_\star from the length n sample path $(X_t)_{t \in [n]}$, and also to construct fully empirical confidence intervals that π_\star and γ_\star with high probability; in particular, the construction

of the intervals should not depend on any unobservable quantities, including π_* and γ_* themselves. As mentioned in the introduction, it is well-known that the *mixing time* of the Markov chain t_{mix} (defined in Eq. 1) is bounded in terms of π_* and γ_* , as shown in Eq. (2). Moreover, convergence rates for empirical processes on Markov chain sequences are also often given in terms of mixing coefficients that can ultimately be bounded in terms of π_* and γ_* (as we will show in the proof of our first result). Therefore, valid confidence intervals for π_* and γ_* can be used to make these rates fully observable.

3 Point estimation

In this section, we present lower and upper bounds on achievable rates for estimating the spectral gap as a function of the length of the sample path n .

3.1 Lower bounds

The purpose of this section is to show lower bounds on the number of observations necessary to achieve a fixed multiplicative (or even just additive) accuracy in estimating the spectral gap γ_* . By Eq. (2), the multiplicative accuracy lower bound for γ_* gives the same lower bound for estimating the mixing time. Our first result holds even for two state Markov chains and shows that a sequence length of $\Omega(1/\pi_*)$ is necessary to achieve even a constant *additive* accuracy in estimating γ_* .

Theorem 1. *Pick any $\bar{\pi} \in (0, 1/4)$. Consider any estimator $\hat{\gamma}_*$ that takes as input a random sample path of length $n \leq 1/(4\bar{\pi})$ from a Markov chain starting from any desired initial state distribution. There exists a two-state ergodic and reversible Markov chain distribution with spectral gap $\gamma_* \geq 1/2$ and minimum stationary probability $\pi_* \geq \bar{\pi}$ such that*

$$\Pr [|\hat{\gamma}_* - \gamma_*| \geq 1/8] \geq 3/8.$$

Next, considering d state chains, we show that a sequence of length $\Omega(d \log(d)/\gamma_*)$ is required to estimate γ_* up to a constant multiplicative accuracy. Essentially, the sequence may have to visit all d states at least $\log(d)/\gamma_*$ times each, on average. This holds *even* if π_* is within a factor of two of the *largest* possible value of $1/d$ that it can take, i.e., when π is nearly uniform.

Theorem 2. *There is an absolute constant $c > 0$ such that the following holds. Pick any positive integer $d \geq 3$ and any $\bar{\gamma} \in (0, 1/2)$. Consider any estimator $\hat{\gamma}_*$ that takes as input a random sample path of length $n < cd \log(d)/\bar{\gamma}$ from a d -state reversible Markov chain starting from any desired initial state distribution. There is an ergodic and reversible Markov chain distribution with spectral gap $\gamma_* \in [\bar{\gamma}, 2\bar{\gamma}]$ and minimum stationary probability $\pi_* \geq 1/(2d)$ such that*

$$\Pr [|\hat{\gamma}_* - \gamma_*| \geq \bar{\gamma}/2] \geq 1/4.$$

The proofs of Theorems 1 and 2 are given in Appendix A.²

3.2 A plug-in based point estimator and its accuracy

Let us now consider the problem of estimating γ_* . For this, we construct a natural plug-in estimator. Along the way, we also provide an estimator for the minimum stationary probability, allowing one to use the bounds from Eq. (2) to trap the mixing time.

Define the random matrix $\widehat{\mathbf{M}} \in [0, 1]^{d \times d}$ and random vector $\hat{\pi} \in \Delta^{d-1}$ by

$$\widehat{M}_{i,j} := \frac{|\{t \in [n-1] : (X_t, X_{t+1}) = (i, j)\}|}{n-1}, \quad i, j \in [d],$$

$$\hat{\pi}_i := \frac{|\{t \in [n] : X_t = i\}|}{n}, \quad i \in [d].$$

Furthermore, define

$$\text{Sym}(\widehat{\mathbf{L}}) := \frac{1}{2}(\widehat{\mathbf{L}} + \widehat{\mathbf{L}}^\top)$$

²A full version of this paper, with appendices, is available on arXiv [31].

to be the symmetrized version of the (possibly non-symmetric) matrix

$$\widehat{\mathbf{L}} := \text{Diag}(\widehat{\boldsymbol{\pi}})^{-1/2} \widehat{\mathbf{M}} \text{Diag}(\widehat{\boldsymbol{\pi}})^{-1/2}.$$

Let $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ be the eigenvalues of $\text{Sym}(\widehat{\mathbf{L}})$. Our estimator of the minimum stationary probability π_* is $\hat{\pi}_* := \min_{i \in [d]} \hat{\pi}_i$, and our estimator of the spectral gap γ_* is $\hat{\gamma}_* := 1 - \max\{\hat{\lambda}_2, |\hat{\lambda}_d|\}$.

These estimators have the following accuracy guarantees:

Theorem 3. *There exists an absolute constant $C > 0$ such that the following holds. Assume the estimators $\hat{\pi}_*$ and $\hat{\gamma}_*$ described above are formed from a sample path of length n from an ergodic and reversible Markov chain. Let $\gamma_* > 0$ denote the spectral gap and $\pi_* > 0$ the minimum stationary probability. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$|\hat{\pi}_* - \pi_*| \leq C \left(\sqrt{\frac{\pi_* \log \frac{d}{\pi_* \delta}}{\gamma_* n}} + \frac{\log \frac{d}{\pi_* \delta}}{\gamma_* n} \right) \quad (4)$$

and

$$|\hat{\gamma}_* - \gamma_*| \leq C \left(\sqrt{\frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\pi_* \delta}}{\pi_* \gamma_* n}} + \frac{\log \frac{1}{\gamma_*}}{\gamma_* n} \right). \quad (5)$$

Theorem 3 implies that the sequence lengths required to estimate π_* and γ_* to within constant multiplicative factors are, respectively, $\tilde{O}\left(\frac{1}{\pi_* \gamma_*}\right)$ and $\tilde{O}\left(\frac{1}{\pi_* \gamma_*^3}\right)$. By Eq. (2), the second of these is also a bound on the required sequence length to estimate t_{mix} .

The proof of Theorem 3 is based on analyzing the convergence of the sample averages $\widehat{\mathbf{M}}$ and $\widehat{\boldsymbol{\pi}}$ to their expectation, and then using perturbation bounds for eigenvalues to derive a bound on the error of $\hat{\gamma}_*$. However, since these averages are formed using a *single sample path* from a (possibly) non-stationary Markov chain, we cannot use standard large deviation bounds; moreover applying Chernoff-type bounds for Markov chains to each entry of $\widehat{\mathbf{M}}$ would result in a significantly worse sequence length requirement, roughly a factor of d larger. Instead, we adapt probability tail bounds for sums of independent random matrices [32] to our non-iid setting by directly applying a blocking technique of [33] as described in the article of [20]. Due to ergodicity, the convergence rate can be bounded without any dependence on the initial state distribution $\boldsymbol{\pi}^{(1)}$. The proof of Theorem 3 is given in Appendix B.

Note that because the eigenvalues of \mathbf{L} are the same as that of the transition probability matrix \mathbf{P} , we could have instead opted to estimate \mathbf{P} , say, using simple frequency estimates obtained from the sample path, and then computing the second largest eigenvalue of this empirical estimate $\widehat{\mathbf{P}}$. In fact, this approach is a way to extend to non-reversible chains, as we would no longer rely on the symmetry of \mathbf{M} or \mathbf{L} . The difficulty with this approach is that \mathbf{P} lacks the structure required by certain strong eigenvalue perturbation results. One could instead invoke the Ostrowski-Elsner theorem [cf. Theorem 1.4 on Page 170 of 34], which bounds the *matching distance* between the eigenvalues of a matrix \mathbf{A} and its perturbation $\mathbf{A} + \mathbf{E}$ by $O(\|\mathbf{E}\|^{1/d})$. Since $\|\widehat{\mathbf{P}} - \mathbf{P}\|$ is expected to be of size $O(n^{-1/2})$, this approach will give a confidence interval for γ_* whose width shrinks at a rate of $O(n^{-1/(2d)})$ —an exponential slow-down compared to the rate from Theorem 3. As demonstrated through an example from [34], the dependence on the d -th root of the norm of the perturbation cannot be avoided in general. Our approach based on estimating a symmetric matrix affords us the use of perturbation results that exploit more structure.

Returning to the question of obtaining a fully empirical confidence interval for γ_* and π_* , we notice that, unfortunately, Theorem 3 falls short of being directly suitable for this, at least without further assumptions. This is because the deviation terms themselves depend inversely both on γ_* and π_* , and hence can never rule out 0 (or an arbitrarily small positive value) as a possibility for γ_* or π_* .³ In effect, the fact that the Markov chain could be slow mixing and the long-term frequency of some

³Using Theorem 3, it is possible to trap γ_* in the union of *two* empirical confidence intervals—one around $\hat{\gamma}_*$ and the other around zero, both of which shrink in width as the sequence length increases.

Algorithm 1 Empirical confidence intervals

Input: Sample path (X_1, X_2, \dots, X_n) , confidence parameter $\delta \in (0, 1)$.

1: Compute state visit counts and smoothed transition probability estimates:

$$N_i := |\{t \in [n-1] : X_t = i\}|, \quad i \in [d];$$
$$N_{i,j} := |\{t \in [n-1] : (X_t, X_{t+1}) = (i, j)\}|, \quad \hat{P}_{i,j} := \frac{N_{i,j} + 1/d}{N_i + 1}, \quad (i, j) \in [d]^2.$$

2: Let $\hat{\mathbf{A}}^\#$ be the group inverse of $\hat{\mathbf{A}} := \mathbf{I} - \hat{\mathbf{P}}$.

3: Let $\hat{\boldsymbol{\pi}} \in \Delta^{d-1}$ be the unique stationary distribution for $\hat{\mathbf{P}}$.

4: Compute eigenvalues $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ of $\text{Sym}(\hat{\mathbf{L}})$, where $\hat{\mathbf{L}} := \text{Diag}(\hat{\boldsymbol{\pi}})^{1/2} \hat{\mathbf{P}} \text{Diag}(\hat{\boldsymbol{\pi}})^{-1/2}$.

5: Spectral gap estimate:

$$\hat{\gamma}_* := 1 - \max\{\hat{\lambda}_2, |\hat{\lambda}_d|\}.$$

6: Empirical bounds for $|\hat{P}_{i,j} - P_{i,j}|$ for $(i, j) \in [d]^2$: $c := 1.01$, $\tau_{n,\delta} := \inf\{t \geq 0 : 2d^2(1 + \lceil \log_c \frac{2n}{t} \rceil_+) e^{-t} \leq \delta\}$,

$$\text{and } \hat{B}_{i,j} := \left(\sqrt{\frac{c\tau_{n,\delta}}{2N_i}} + \sqrt{\frac{c\tau_{n,\delta}}{2N_i} + \sqrt{\frac{2c\hat{P}_{i,j}(1 - \hat{P}_{i,j})\tau_{n,\delta}}{N_i} + \frac{(5/3)\tau_{n,\delta} + |\hat{P}_{i,j} - 1/d|}{N_i}}} \right)^2.$$

7: Relative sensitivity of $\boldsymbol{\pi}$:

$$\hat{\kappa} := \frac{1}{2} \max \left\{ \hat{A}_{j,j}^\# - \min \left\{ \hat{A}_{i,j}^\# : i \in [d] \right\} : j \in [d] \right\}.$$

8: Empirical bounds for $\max_{i \in [d]} |\hat{\pi}_i - \pi_i|$ and $\max \bigcup_{i \in [d]} \{|\sqrt{\pi_i/\hat{\pi}_i} - 1|, |\sqrt{\hat{\pi}_i/\pi_i} - 1|\}$:

$$\hat{b} := \hat{\kappa} \max \left\{ \hat{B}_{i,j} : (i, j) \in [d]^2 \right\}, \quad \hat{\rho} := \frac{1}{2} \max \bigcup_{i \in [d]} \left\{ \frac{\hat{b}}{\hat{\pi}_i}, \frac{\hat{b}}{[\hat{\pi}_i - \hat{b}]_+} \right\}.$$

9: Empirical bounds for $|\hat{\gamma}_* - \gamma_*|$:

$$\hat{w} := 2\hat{\rho} + \hat{\rho}^2 + (1 + 2\hat{\rho} + \hat{\rho}^2) \left(\sum_{(i,j) \in [d]^2} \frac{\hat{\pi}_i}{\hat{\pi}_j} \hat{B}_{i,j}^2 \right)^{1/2}.$$

states could be small makes it difficult to be confident in the estimates provided by $\hat{\gamma}_*$ and $\hat{\pi}_*$. This suggests that in order to obtain fully empirical confidence intervals, we need an estimator that is not subject to such effects—we pursue this in Section 4. Theorem 3 thus primarily serves as a point of comparison for what is achievable in terms of estimation accuracy when one does not need to provide empirical confidence bounds.

4 Fully empirical confidence intervals

In this section, we address the shortcoming of Theorem 3 and give fully empirical confidence intervals for the stationary probabilities and the spectral gap γ_* . The main idea is to use the Markov property to eliminate the dependence of the confidence intervals on the unknown quantities (including π_* and γ_*). Specifically, we estimate the transition probabilities from the sample path using simple frequency estimates: as a consequence of the Markov property, for each state, the frequency estimates converge at a rate that depends only on the number of visits to the state, and in particular the rate (given the visit count of the state) is independent of the mixing time of the chain.

As discussed in Section 3, it is possible to form a confidence interval for γ_* based on the eigenvalues of an estimated transition probability matrix by appealing to the Ostrowski-Elsner theorem. However, as explained earlier, this would lead to a slow $O(n^{-1/(2d)})$ rate. We avoid this slow rate by using an estimate of the symmetric matrix L , so that we can use a stronger perturbation result (namely Weyl's inequality, as in the proof of Theorem 3) available for symmetric matrices.

To form an estimate of L based on an estimate of the transition probabilities, one possibility is to estimate π using a frequency-based estimate for π as was done in Section 3, and appeal to the relation $L = \text{Diag}(\pi)^{1/2} P \text{Diag}(\pi)^{-1/2}$ to form a plug-in estimate. However, as noted in Section 3.2, confidence intervals for the entries of π formed this way may depend on the mixing time. Indeed, such an estimate of π does not exploit the Markov property.

We adopt a different strategy for estimating π , which leads to our construction of empirical confidence intervals, detailed in Algorithm 1. We form the matrix \hat{P} using smoothed frequency estimates of P (Step 1), then compute the so-called group inverse $\hat{A}^\#$ of $\hat{A} = I - \hat{P}$ (Step 2), followed by finding the unique stationary distribution $\hat{\pi}$ of \hat{P} (Step 3), this way decoupling the bound on the accuracy of $\hat{\pi}$ from the mixing time. The group inverse $\hat{A}^\#$ of \hat{A} is uniquely defined; and if \hat{P} defines an ergodic chain (which is the case here due to the use of the smoothed estimates), $\hat{A}^\#$ can be computed at the cost of inverting an $(d-1) \times (d-1)$ matrix [35, Theorem 5.2].⁴ Further, once given $\hat{A}^\#$, the unique stationary distribution $\hat{\pi}$ of \hat{P} can be read out from the last row of $\hat{A}^\#$ [35, Theorem 5.3]. The group inverse is also used to compute the sensitivity of π . Based on $\hat{\pi}$ and \hat{P} , we construct the plug-in estimate \hat{L} of L , and use the eigenvalues of its symmetrization to form the estimate $\hat{\gamma}_*$ of the spectral gap (Steps 4 and 5). In the remaining steps, we use perturbation analyses to relate $\hat{\pi}$ and π , viewing P as the perturbation of \hat{P} ; and also to relate $\hat{\gamma}_*$ and γ_* , viewing L as a perturbation of $\text{Sym}(\hat{L})$. Both analyses give error bounds entirely in terms of observable quantities (e.g., $\hat{\kappa}$), tracing back to empirical error bounds for the smoothed frequency estimates of P .

The most computationally expensive step in Algorithm 1 is the computation of the group inverse $\hat{A}^\#$, which, as noted reduces to matrix inversion. Thus, with a standard implementation of matrix inversion, the algorithm's time complexity is $O(n + d^3)$, while its space complexity is $O(d^2)$.

To state our main theorem concerning Algorithm 1, we first define κ to be analogous to $\hat{\kappa}$ from Step 7, with $\hat{A}^\#$ replaced by the group inverse $A^\#$ of $A := I - P$. The result is as follows.

Theorem 4. *Suppose Algorithm 1 is given as input a sample path of length n from an ergodic and reversible Markov chain and confidence parameter $\delta \in (0, 1)$. Let $\gamma_* > 0$ denote the spectral gap, π the unique stationary distribution, and $\pi_* > 0$ the minimum stationary probability. Then, on an event of probability at least $1 - \delta$,*

$$\pi_i \in [\hat{\pi}_i - \hat{b}, \hat{\pi}_i + \hat{b}] \quad \text{for all } i \in [d], \quad \text{and} \quad \gamma_* \in [\hat{\gamma}_* - \hat{w}, \hat{\gamma}_* + \hat{w}].$$

Moreover, \hat{b} and \hat{w} almost surely satisfy (as $n \rightarrow \infty$)

$$\hat{b} = O\left(\max_{(i,j) \in [d]^2} \kappa \sqrt{\frac{P_{i,j} \log \log n}{\pi_i n}}\right), \quad \hat{w} = O\left(\frac{\kappa}{\pi_*} \sqrt{\frac{\log \log n}{\pi_* n}} + \sqrt{\frac{d \log \log n}{\pi_* n}}\right).⁵$$

The proof of Theorem 4 is given in Appendix C. As mentioned above, the obstacle encountered in Theorem 3 is avoided by exploiting the Markov property. We establish fully observable upper and lower bounds on the entries of P that converge at a $\sqrt{n/\log \log n}$ rate using standard martingale tail inequalities; this justifies the validity of the bounds from Step 6. Properties of the group inverse [35, 36] and eigenvalue perturbation theory [34] are used to validate the empirical bounds on π_i and γ_* developed in the remaining steps of the algorithm.

The first part of Theorem 4 provides valid empirical confidence intervals for each π_i and for γ_* , which are simultaneously valid at confidence level δ . The second part of Theorem 4 shows that the

⁴ The group inverse of a square matrix A , a special case of the Drazin inverse, is the unique matrix $A^\#$ satisfying $AA^\#A = A$, $A^\#AA^\# = A^\#$ and $A^\#A = AA^\#$.

⁵In Theorems 4 and 5, our use of big- O notation is as follows. For a random sequence $(Y_n)_n$ and a (non-random) positive sequence $(\varepsilon_{\theta,n})_n$ parameterized by θ , we say " $Y_n = O(\varepsilon_{\theta,n})$ holds almost surely as $n \rightarrow \infty$ " if there is some universal constant $C > 0$ such that for all θ , $\limsup_{n \rightarrow \infty} Y_n/\varepsilon_{\theta,n} \leq C$ holds almost surely.

width of the intervals decrease as the sequence length increases. We show in Appendix C.5 that $\kappa \leq d/\gamma_*$, and hence $\hat{b} = O\left(\max_{(i,j) \in [d]^2} \frac{d}{\gamma_*} \sqrt{\frac{P_{i,j} \log \log n}{\pi_i n}}\right)$, $\hat{w} = O\left(\frac{d}{\pi_* \gamma_*} \sqrt{\frac{\log \log n}{\pi_* n}}\right)$.

It is easy to combine Theorems 3 and 4 to yield intervals whose widths shrink at least as fast as both the non-empirical intervals from Theorem 3 and the empirical intervals from Theorem 4. Specifically, determine lower bounds on π_* and γ_* using Algorithm 1, $\pi_* \geq \min_{i \in [d]} [\hat{\pi}_i - \hat{b}]_+$, $\gamma_* \geq [\hat{\gamma}_* - \hat{w}]_+$; then plug-in these lower bounds for π_* and γ_* in the deviation bounds in Eq. (5) from Theorem 3. This yields a new interval centered around the estimate of γ_* from Theorem 3, and it no longer depends on unknown quantities. The interval is a valid $1 - 2\delta$ probability confidence interval for γ_* , and for sufficiently large n , the width shrinks at the rate given in Eq. (5). We can similarly construct an empirical confidence interval for π_* using Eq. (4), which is valid on the same $1 - 2\delta$ probability event.⁶ Finally, we can take the intersection of these new intervals with the corresponding intervals from Algorithm 1. This is summarized in the following theorem, which we prove in Appendix D.

Theorem 5. *The following holds under the same conditions as Theorem 4. For any $\delta \in (0, 1)$, the confidence intervals \hat{U} and \hat{V} described above for π_* and γ_* , respectively, satisfy $\pi_* \in \hat{U}$ and $\gamma_* \in \hat{V}$ with probability at least $1 - 2\delta$. Furthermore, the widths of these intervals almost surely satisfy (as $n \rightarrow \infty$) $|\hat{U}| = O\left(\sqrt{\frac{\pi_* \log \frac{d}{\pi_* \delta}}{\gamma_* n}}\right)$, $|\hat{V}| = O\left(\min\left\{\sqrt{\frac{\log \frac{d}{\delta} \cdot \log(n)}{\pi_* \gamma_* n}}, \hat{w}\right\}\right)$, where \hat{w} is the width from Algorithm 1.*

5 Discussion

The construction used in Theorem 5 applies more generally: Given a confidence interval of the form $I_n = I_n(\gamma_*, \pi_*, \delta)$ for some confidence level δ and a fully empirical confidence set $E_n(\delta)$ for (γ_*, π_*) for the same level, $I'_n = E_n(\delta) \cap \cup_{(\gamma, \pi) \in E_n(\delta)} I_n(\gamma, \pi, \delta)$ is a valid fully empirical 2δ -level confidence interval whose asymptotic width matches that of I_n up to lower order terms under reasonable assumptions on E_n and I_n . In particular, this suggests that future work should focus on closing the gap between the lower and upper bounds on the accuracy of point-estimation. Another interesting direction is to reduce the computation cost: The current cubic cost in the number of states can be too high even when the number of states is only moderately large.

Perhaps more important, however, is to extend our results to large state space Markov chains: In most practical applications the state space is continuous or is exponentially large in some natural parameters. As follows from our lower bounds, without further assumptions, the problem of fully data dependent estimation of the mixing time is intractable for information theoretical reasons. Interesting directions for future work thus must consider Markov chains with specific structure. Parametric classes of Markov chains, including but not limited to Markov chains with factored transition kernels with a few factors, are a promising candidate for such future investigations. The results presented here are a first step in the ambitious research agenda outlined above, and we hope that they will serve as a point of departure for further insights in the area of fully empirical estimation of Markov chain parameters based on a single sample path.

References

- [1] J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. Springer-Verlag, 2001.
- [2] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning)*. A Bradford Book, 1998.
- [3] D. Levin, Y. Peres, and E. Wilmer. *Markov Chains and Mixing Times*. AMS, 2008.
- [4] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer, 1993.
- [5] C. Kipnis and S. R. S. Varadhan. Central limit theorem for additive functionals of reversible markov processes and applications to simple exclusions. *Comm. Math. Phys.*, 104(1):1–19, 1986.

⁶For the π_* interval, we only plug-in lower bounds on π_* and γ_* only where these quantities appear as $1/\pi_*$ and $1/\gamma_*$ in Eq. (4). It is then possible to “solve” for observable bounds on π_* . See Appendix D for details.

- [6] I. Kontoyiannis, L. A. Lastras-Montaño, and S. P. Meyn. Exponential bounds and stopping rules for MCMC and general Markov chains. In *VALUETOOLS*, page 45, 2006.
- [7] M.-F. Balcan, A. Beygelzimer, and J. Langford. Agnostic active learning. In *ICML*, pages 65–72, 2006.
- [8] V. Mnih, Cs. Szepesvári, and J.-Y. Audibert. Empirical Bernstein stopping. In *ICML*, pages 672–679, 2008.
- [9] A. Maurer and M. Pontil. Empirical Bernstein bounds and sample-variance penalization. In *COLT*, 2009.
- [10] L. Li, M. L. Littman, T. J. Walsh, and A. L. Strehl. Knows what it knows: a framework for self-aware learning. *Machine Learning*, 82(3):399–443, 2011.
- [11] J. M. Flegal and G. L. Jones. Implementing MCMC: estimating with confidence. In *Handbook of Markov chain Monte Carlo*, pages 175–197. Chapman & Hall/CRC, 2011.
- [12] B. M. Gyori and D. Paulin. Non-asymptotic confidence intervals for MCMC in practice. arXiv:1212.2016, 2014.
- [13] A. Swaminathan and T. Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *ICML*, 2015.
- [14] D. Gillman. A Chernoff bound for random walks on expander graphs. *SIAM Journal on Computing*, 27(4):1203–1220, 1998.
- [15] C. A. León and F. Perron. Optimal Hoeffding bounds for discrete reversible Markov chains. *Annals of Applied Probability*, pages 958–970, 2004.
- [16] D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20:1–32, 2015.
- [17] S. T. Garren and R. L. Smith. Estimating the second largest eigenvalue of a Markov transition matrix. *Bernoulli*, 6:215–242, 2000.
- [18] G. L. Jones and J. P. Hobert. Honest exploration of intractable probability distributions via markov chain monte carlo. *Statist. Sci.*, 16(4):312–334, 11 2001.
- [19] Y. Atchadé. Markov Chain Monte Carlo confidence intervals. *Bernoulli*, 2015. (to appear).
- [20] B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. *The Annals of Probability*, 22(1):94–116, January 1994.
- [21] R. L. Karandikar and M. Vidyasagar. Rates of uniform convergence of empirical means with mixing processes. *Statistics and Probability Letters*, 58(3):297–307, 2002.
- [22] D. Gamarnik. Extension of the PAC framework to finite and countable Markov chains. *IEEE Transactions on Information Theory*, 49(1):338–345, 2003.
- [23] M. Mohri and A. Rostamizadeh. Stability bounds for non-iid processes. In *NIPS*, 2008.
- [24] M. Mohri and A. Rostamizadeh. Rademacher complexity bounds for non-i.i.d. processes. In *NIPS*, 2009.
- [25] I. Steinwart and A. Christmann. Fast learning from non-i.i.d. observations. In *NIPS*, 2009.
- [26] I. Steinwart, D. Hush, and C. Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100(1):175–194, 2009.
- [27] D. McDonald, C. Shalizi, and M. Schervish. Estimating beta-mixing coefficients. In *AISTATS*, pages 516–524, 2011.
- [28] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *FOCS*, pages 259–269. IEEE, 2000.
- [29] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing closeness of discrete distributions. *Journal of the ACM (JACM)*, 60(1):4:2–4:25, 2013.
- [30] N. Bhatnagar, A. Bogdanov, and E. Mossel. The computational complexity of estimating MCMC convergence time. In *RANDOM*, pages 424–435. Springer, 2011.
- [31] D. Hsu, A. Kontorovich, and C. Szepesvári. Mixing time estimation in reversible Markov chains from a single sample path. *CoRR*, abs/1506.02903, 2015.
- [32] J. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 2015.
- [33] S. Bernstein. Sur l’extension du theoreme limite du calcul des probabilités aux sommes de quantités dependantes. *Mathematische Annalen*, 97:1–59, 1927.
- [34] G. W. Stewart and J. Sun. *Matrix perturbation theory*. Academic Press, Boston, 1990.
- [35] C. D. Meyer Jr. The role of the group generalized inverse in the theory of finite Markov chains. *SIAM Review*, 17(3):443–464, 1975.
- [36] G. Cho and C. Meyer. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications*, 335:137–150, 2001.