Statistical estimation of ergodic Markov chain kernel over discrete state space

GEOFFREY WOLFER[*] and ARYEH KONTOROVICH[†]

Computer Science Department, Ben-Gurion University of the Negev, Beer-Sheva, Israel.
E-mail: [*] *geoffrey@post.bgu.ac.il;* [†] *karyeh@cs.bgu.ac.il*

We investigate the statistical complexity of estimating the parameters of a discrete-state Markov chain kernel from a single long sequence of state observations. In the finite case, we characterize (modulo logarithmic factors) the minimax sample complexity of estimation with respect to the operator infinity norm, while in the countably infinite case, we analyze the problem with respect to a natural entry-wise norm derived from total variation. We show that in both cases, the sample complexity is governed by the mixing properties of the unknown chain, for which, in the finite-state case, there are known finite-sample estimators with fully empirical confidence intervals.

Keywords: discrete state space; ergodic Markov chain; minimax theory

1. Introduction

Approximately recovering the parameters of a discrete distribution is a classical problem in computer science and statistics (see, e.g., Han et al. [9], Kamath et al. [13], Orlitsky and Suresh [21] and the references therein). Total variation (TV) is a natural and well-motivated choice of approximation metric (Devroye and Lugosi [7]), and the two metrics we use throughout the paper will be derived from TV. The minimax sample complexity for obtaining an ε-approximation to the unknown distribution in TV (but see Waggoner [27] for results on other ℓ_p norms) is well known to be of the order of d/ε^2, where d is the support size (see, e.g., Anthony and Bartlett [1], Kontorovich and Pinelis [16]).

This paper deals with estimating the transition probability kernel of a discrete state time-homogeneous Markov chain in the minimax setting. The Markov case is much less well-understood than the i.i.d. one. The main additional complexity introduced by the Markov case on top of the i.i.d. one is that the sample complexity involves not only the number of states and the precision parameter ε, but also the chain's mixing properties.

Our contributions

In the finite-state case, we compute, up to logarithmic factors, (apparently the first, in *any* metric) high-probability minimax sample complexity for the estimation problem in the Markovian setting, which seeks to recover, from a single long run of an unknown Markov chain, the values of its transition matrix up to a tolerance of ε in the $\| \cdot \|_\infty$ operator norm. We obtain upper and lower bounds on the sample complexity (sequence length) in terms of ε, the number of states, the stationary distribution, and mixing time of the Markov chain.

In the countably infinite case, for a natural class of chains and with respect to an entry-wise metric derived from TV, we derive an upper bound on the sample complexity that depends in a delicate way on some measure of complexity of the kernel, precision ε and mixing time, and provide sufficient conditions on the kernel and initial distribution to obtain convergence guarantees.

2. Definitions and notation

We denote by Ω the state space of the Markov chain and by m the size of the sample received by the estimation procedure. The simplex of all distributions over Ω will be denoted by Δ_Ω, and the set of all Markov kernels by $\mathcal{M}_\Omega$. For $|\Omega| < \infty$, we put $d := |\Omega|$ and $[d] = \Omega = \{1, 2, \ldots, d\}$. For $\boldsymbol{\mu} \in \Delta_\Omega$, we will write either $\boldsymbol{\mu}(i)$ or μ_i, as dictated by esthetics and convenience. All vectors are rows unless indicated otherwise. We use the standard total variation norm, which, up to a convention-dependent factor of 2, corresponds to the ℓ_1 norm: $2\|\boldsymbol{x}\|_{\mathsf{TV}} = \|\boldsymbol{x}\|_1 = \sum_{i \in \Omega} |x_i|$. We assume familiarity with basic Markov chain concepts (see, e.g., Kemeny and Snell [15], Levin et al. [19]). A time-homogeneous Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$ on state space Ω is specified by an initial distribution $\boldsymbol{\mu} \in \Delta_\Omega$ and a kernel $\boldsymbol{M} \in \mathcal{M}_\Omega$ in the usual way: $(X_1, \ldots, X_m) \sim (\boldsymbol{M}, \boldsymbol{\mu})$ means that

$$\mathbb{P}\big((X_1, \ldots, X_m) = (x_1, \ldots, x_m)\big) = \boldsymbol{\mu}(x_1) \prod_{t=1}^{m-1} \boldsymbol{M}(x_t, x_{t+1}).$$

We write $\mathbb{P}_{\boldsymbol{M}, \mu}(\cdot)$ to denote probabilities over sequences induced by the Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$, and omit one or both subscripts when clear from context. We say that $\boldsymbol{\pi}$ is a *stationary distribution* for $\boldsymbol{M}$ if $\boldsymbol{\pi} \boldsymbol{M} = \boldsymbol{\pi}$, and that the Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$ is *stationary* if $\boldsymbol{\mu} = \boldsymbol{\pi}$. We will assume the chain to be *irreducible* and *positive recurrent*. Namely, $\boldsymbol{M}$ consists of a single communicating class, and defining the *return time* of state i as $T_i = \min\{t \geq 0 : X_t = i\}$, we have that for any state $i \in \Omega$, $\mathbb{E}[T_i] < \infty$. This is sufficient to guarantee existence of a stationary $\boldsymbol{\pi}$. We will further restrict our analysis to *geometrically ergodic* Markov chains to enable spectral methods.

Definition 2.1 (Geometric ergodicity, Roberts et al. [23]). The chain $(\boldsymbol{M}, \boldsymbol{\mu})$ with stationary distribution $\boldsymbol{\pi}$ is geometrically ergodic if there is a $\rho \in (0, 1)$ and for all $i \in \Omega$ there is a $C_i \in \mathbb{R}_+$ such that

$$\big\| \boldsymbol{M}^t(i, \cdot) - \boldsymbol{\pi} \big\|_{\mathsf{TV}} \leq C_i \rho^t, \quad t \in \mathbb{N}.$$

Any chain that satisfies all the above properties will henceforth simply be called *ergodic*, and all chains mentioned in this work will be assumed ergodic unless stated otherwise. If $\boldsymbol{M}$ is ergodic with stationary distribution $\boldsymbol{\pi}$, then $\boldsymbol{\pi}$ is necessarily unique. To any Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$, we associate the following measure of non-stationarity

$$\|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}}^2 \doteq \sum_{i \in \Omega} \boldsymbol{\mu}(i)^2/\boldsymbol{\pi}(i) \in [1, \infty], \tag{2.1}$$

where the $\|\cdot\|_{2,\boldsymbol{\pi}}$ norm is induced by the inner product in the Hilbert space $\ell_2(\boldsymbol{\pi})$ (Levin et al. [19], Chapter 12). When $|\Omega| < \infty$, we can define the *minimum stationary probability* by

$$\pi_\star \doteq \min_{i \in \Omega} \boldsymbol{\pi}(i). \tag{2.2}$$

In this case, by ergodicity $\pi_\star > 0$ and $\|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}}^2 \leq \frac{1}{\pi_\star} < \infty$. The *mixing time* of an ergodic $\boldsymbol{M}$ is defined by

$$t_{\mathrm{mix}} \doteq \inf\left\{ t \geq 1 : \sup_{\mu \in \Delta_\Omega} \big\| \boldsymbol{\mu} \boldsymbol{M}^t - \boldsymbol{\pi} \big\|_{\mathsf{TV}} < \frac{1}{4} \right\}. \tag{2.3}$$

We define $\boldsymbol{Q} \doteq \mathrm{diag}(\boldsymbol{\pi}) \boldsymbol{M}$ as the matrix $\boldsymbol{Q}(i, j) = \mathbb{P}_{\boldsymbol{\pi}}(X_t = i, X_{t+1} = j)$. A chain $\boldsymbol{M} \in \mathcal{M}_\Omega$ is said to be *reversible* if $\boldsymbol{Q}^\mathsf{T} = \boldsymbol{Q}$. The eigenvalues of an ergodic and reversible $\boldsymbol{M}$ lie in $(-1, 1]$, and thus

may be ordered (counting multiplicities): $1 = \lambda_1 > \lambda_2 \geq \cdots \geq \lambda_d > -1$. The *spectral gap* and *absolute spectral gap* of a reversible chain are defined, respectively, by

$$\gamma \doteq 1 - \lambda_2 \quad \text{and} \quad \gamma_\star \doteq 1 - \max\{\lambda_2, |\lambda_d|\}. \tag{2.4}$$

Paulin [22] generalizes the *multiplicative reversiblization* approach of Fill [8] by defining the *pseudo-spectral gap*

$$\gamma_{\mathsf{ps}} \doteq \max_{k \geq 1}\{\gamma\big((M^\dagger)^k M^k\big)/k\}, \tag{2.5}$$

where $M^\dagger$ is the *time reversal* of M – the adjoint of M under $\ell_2(\pi)$ – given by $M^\dagger(i, j) \doteq \pi(j)M(j, i)/\pi(i)$.

For a linear operator $A : \Omega \to \Omega$,

$$\|A\|_\infty = \sup_{i \in \Omega} \sum_{j \in \Omega} |A(i, j)| \tag{2.6}$$

is the operator norm induced by ℓ_∞ (Horn and Johnson [11]). We also define the following entry-wise norm

$$\||A\|| \doteq \sum_{(i,j) \in \Omega^2} |A(i, j)|. \tag{2.7}$$

The norms in (2.6) and (2.7) induce our two notions of distance between Markov kernels M, M' with respective stationary distributions π and π':

$$\|M - M'\|_\infty \quad \text{and} \quad \||Q - Q'\||.$$

For any $M \in \mathcal{M}_\Omega$, define its *Dobrushin contraction* coefficient

$$\kappa(M) \doteq \max_{(i,j) \in \Omega^2} \|M(i, \cdot) - M(j, \cdot)\|_{\mathsf{TV}}; \tag{2.8}$$

this quantity is also associated with Döblin's name. The term "contraction" refers to the property

$$\|(\mu - \mu')M\|_{\mathsf{TV}} \leq \kappa(M)\|\mu - \mu'\|_{\mathsf{TV}}, \quad (\mu, \mu') \in \Delta_\Omega^2, \tag{2.9}$$

which was observed by Markov [20], §5 (see Kontorovich and Ramanan [18], Lemma A.2, for an elementary proof).

3. Main results

In Section 3.1, we formally state the minimax results for the finite state setting, and then exhibit our results for the countably infinite case in Section 3.2.

3.1. Estimation with respect to $\| \cdot \|_\infty$ for finite Ω

Theorem 3.1 (Sample complexity upper bound w.r.t $\| \cdot \|_\infty$ when $|\Omega| < \infty$). *Let* $\varepsilon \in (0, 2)$, $\delta \in (0, 1)$, *and let* $X = (X_1, \ldots, X_m) \sim (M, \mu)$, M *ergodic with stationary distribution* π. *Then an*

estimator $\widehat{\boldsymbol{M}} : \Omega^m \to \mathcal{M}_\Omega$ *exists such that whenever*

$$m \geq c \max \left\{ \frac{1}{\varepsilon^2 \pi_\star} \max \left\{ d, \ln \frac{1}{\varepsilon \delta} \right\}, \frac{1}{\gamma_{\mathsf{ps}} \pi_\star} \ln \frac{d \|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}}}{\delta} \right\}$$

we have, with probability at least $1 - \delta$,

$$\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_\infty < \varepsilon,$$

where c is a universal constant, $d = |\Omega|$, γ_{ps} *is the* pseudo-spectral gap (2.5), $\pi_\star$ *the minimum station-ary probability* (2.2), *and* $\|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}}^2 \leq 1/\pi_\star$ *is defined in* (2.1).

Although the sample complexity depends on the spectral quantity γ_{ps}, and minimal stationary prob-ability $\pi_\star$ of the unknown chain, these can be efficiently estimated with finite-sample data-dependent confidence intervals from a single trajectory ([Hsu et al. [12], Wolfer and Kontorovich [31]). Moreover, even though the upper bound formally depends on the unknown (and, in our one-trajectory setting, im-possible to estimate) initial distribution $\boldsymbol{\mu}$, we note that (i) this dependence is only logarithmic and (ii) an upper bound on $\|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}}^2$ in terms of $\pi_\star$ is easily provided.

Remark 3.1. This upper bound is superior to the one given at Wolfer and Kontorovich [30], Theo-rem 1, shaving a multiplicative factor of $\ln d$ off the first term, except in the extremely high precision regime where $\ln \frac{1}{\varepsilon} \geq d \ln d$.

Theorem 3.2 (Sample complexity lower bound w.r.t $\|\cdot\|_\infty$ **when** $|\Omega| < \infty$**).** *For every* $\varepsilon \in (0, 1/32)$, $\gamma_{\mathsf{ps}} \in (0, 1/8)$, $d = 6k \geq 12$, *and every estimation procedure, there exists a d-state Markov chain* $\boldsymbol{M}$ *with pseudo-spectral gap* γ_{ps} *and stationary distribution* $\boldsymbol{\pi}$ *such that the estimation procedure must require a sequence* $\boldsymbol{X} = (X_1, \ldots, X_m)$ *drawn from the unknown* $\boldsymbol{M}$ *of length at least*

$$m \geq c \max \left\{ \frac{d}{\varepsilon^2 \pi_\star}, \frac{d \ln d}{\gamma_{\mathsf{ps}}} \right\},$$

where c is a universal constant, to ensure $\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_\infty < \varepsilon$ *with probability greater than* 9/10, *and where* d, γ_{ps}, $\pi_\star$ *are as in Theorem* 3.1.

The proof of Theorem 3.2 actually yields a bit more than claimed in the statement. For any $\pi_\star \in (0, 1/d]$, a Markov chain $\boldsymbol{M}$ can be constructed that achieves the $\frac{d}{\varepsilon^2 \pi_\star}$ component of the bound. Additionally, the $\frac{d}{\gamma_{\mathsf{ps}}}$ component is achievable by a class of *reversible* Markov chains with spectral gap $\gamma \leq \gamma_{\mathsf{ps}} \leq 2\gamma$, and uniform stationary distribution – for which $\pi_\star = 1/d$ – exhibiting tightness of the obtained bound.

The form of the lower bound indicates that in some regimes, estimating the pseudo-spectral gap up to constant multiplicative error, which requires $\frac{d}{\gamma_{\mathsf{ps}}}$ (Hsu et al. [12], Wolfer and Kontorovich [31]), is as difficult as estimating the entire transition matrix (for our choice of metric $\|\cdot\|_\infty$). We stress that our procedure and guarantees only require ergodicity (and not, say, reversibility) to work.

3.2. Results for estimation with respect to $\|\|\cdot\|\|$

Over an infinite space, $\pi_\star = 0$ conveys no information, which motivates an alternative notion of dis-tance. For a chain $\boldsymbol{M}$, the kernel of *doublet frequencies* $\boldsymbol{Q} = \mathrm{diag}(\boldsymbol{\pi})\boldsymbol{M}$ encodes all information about

an ergodic chain (Vidyasagar [26]), and for two such operators, Q and Q' it is the case that

$$Q = Q' \implies M = M'.$$

Further, it is easily verified that

$$\|\| Q - Q' \|\| = 2\| Q - Q' \|_{\mathrm{TV}},$$

where we see Q and Q' as distributions over $\Omega \times \Omega$.

Remark 3.2. The loss of our estimation problem is distinct from the one considered in Hao et al. [10], which weights the state-wise expected loss with respect to the stationary distribution of the chain, and also allows for sample bounds independent of $\pi_\star$.

Theorem 3.3. *Let $\varepsilon \in (0, 2)$, $\delta \in (0, 1)$, and $X = (X_1, \ldots, X_m) \sim (M, \mu)$, M ergodic with stationary distribution π, and write $Q = \mathrm{diag}(\pi)M$. There exists an estimator $\widehat{Q} : \Omega^m \to \Delta_{\Omega \times \Omega}$ such that for*

$$m \geq c \frac{t_{\mathrm{mix}}}{\varepsilon^2} \max\left\{ \|\| Q \|\|_{1/2}, \ln\left(\frac{\|\mu/\pi\|_{2,\pi}}{\delta} \right) \right\},$$

we have $\|\| \widehat{Q} - Q \|\| < \varepsilon$ with probability at least $1 - \delta$, where

$$\|\| Q \|\|_{1/2} \doteq \left(\sum_{(i,j) \in \Omega^2} \sqrt{Q(i,j)} \right)^2,$$

c is a universal constant, t_{mix} is defined at (2.3), and $\|\mu/\pi\|_{2,\pi}$ is defined in (2.1).

Remark 3.3. Necessary conditions for the upper bound to be non-vacuous are that both $\|\| Q \|\|_{1/2} < \infty$ and $\|\mu/\pi\|_{2,\pi} < \infty$. Importantly, $\|\mu/\pi\|_{2,\pi} < \infty$ implies but is not implied by $\mu \ll \pi$; take, for example, $\mu(i) \propto \frac{1}{i^2}$ and $\pi(i) \propto \frac{1}{i^4}$. Notice that in the special case where $|\Omega| = d < \infty$, we have $\|\| Q \|\|_{1/2} \leq d^2$ and the bound reduces to $t_{\mathrm{mix}} \frac{d^2}{\varepsilon^2}$ (up to logarithmic factors). The mixing time t_{mix}, unknown a priori, can be estimated with finite-sample empirical intervals (Wolfer [29]).

4. Overview of techniques

4.1. Estimating with respect to the $\| \cdot \|_\infty$ norm when $|\Omega| < \infty$

The upper bound for the estimation problem in Theorem 3.1 is achieved by a (mildly smoothed) natural estimator defined at the beginning of Section 6.1. If the stationary distribution is bounded away from 0, the chain will visit each state a constant fraction of the total sequence length. Exponential concentration (controlled by the spectral gap) provides high-probability confidence intervals about the expectations. A technical complication is that the empirical distribution of the transitions out of a state i, conditional on the number of visits N_i to that state, is not binomial but actually rather complicated – this is due to the fact that the sequence length is fixed and so a large value of N_i "crowds out" other observations. We overcome this by simulating a trajectory from the Markov chain with an array of independent random variables, as described in Billingsley [4], p. 19. The factor $\|\mu/\pi\|_{2,\pi}$ in the bounds quantifies the price one pays for not assuming (as we do not) stationarity of the unknown Markov chain.

Our chief technical contribution is in establishing the sample complexity lower bounds for the finite space estimation problem. We do this by constructing two distinct lower bounds.

The lower bound of $\frac{d \ln d}{\gamma_{\mathrm{ps}}}$ is derived by a covering argument and a classical reduction scheme to a collection of testing problems using a class of reversible Markov chains we construct, with a carefully controlled pseudo-spectral gap.[1] The latter can be upper and lower bounded up to universal constants in three key steps. First, we leverage the block structure of the transition matrix of the non-perturbed member of the family to compute its entire spectrum explicitly (Lemma 6.5), and deduce its absolute spectral gap. We then extend the bound to other members of the family, using *Markov chain comparison techniques*, going through a well known variational definition of the spectral gap. Finally, we conclude by showing that the pseudo-spectral and spectral gap are within a factor of 2 for our class of symmetric Markov chains.

The lower bound of $\frac{d}{\varepsilon^2 \pi_\star}$ is based on the observation that estimating the whole kernel is at least as hard as estimating the conditional distribution a single state. From here, we construct a class of matrices where one state is both hard to reach and difficult to estimate, by constructing mixture of indistinguishable distributions for that particular state, indexed by a large subset of the binary hypercube. We express the statistical distance between words of length m distributed according to different matrices of this class in terms of $\pi_\star$ and the KL divergence between the conditional distributions of the hard-to-reach state, by taking advantage of the structure of the class, and invoke an argument from Tsybakov to conclude ours.

4.2. Estimating with respect to the $\|\|\cdot\|\|$ metric

The extension to a countably infinite setting requires an alternative notion of distance between chains. The proof then introduces the natural counting estimator $\widehat{Q}(i, j) = \frac{N_{ij}}{m-1}$ of transitions from i to j, and starts by controlling the error in expectation. It reduces the problem to the study of the variance of the random variable N_{ij}, which is achieved by constructing another Markov chain with approximately the same mixing time, and invoking known results from Paulin [22] for the variance of sums of functions under the Markovian setting. The result is then obtained by controlling the fluctuations around this expectation by a bounded differences argument.

5. Related work

Our Markov chain statistical estimation setup is a natural extension of the PAC distribution learning model of Kearns et al. [14]. Despite the plethora of literature on estimating Markov transition matrices, (see, e.g., Billingsley [4], Craig and Sendi [6], Welton and Ades [28]) we were not able to locate any rigorous finite-sample PAC-type results.

The minimax problem has recently received some attention, and Hao et al. [10] have, in parallel to us, shown the first minimax bounds, in expectation, for the problem of estimating the transition matrix M of a Markov chain under a certain class of divergences. The authors consider the case where $\min_{i,j} M(i, j) \geq \alpha > 0$, essentially showing that for some family of smooth f-divergences, the expected risk is of the order of $\frac{d f''(1)}{m \pi_\star}$. The metric used in this paper is based on TV, which corresponds to the f-divergence induced by $f(t) = \frac{1}{2}|t - 1|$, which is not differentiable at $t = 1$. The results of Hao

---

[1]The family of chains used in the lower bound of Hsu et al. [12] does not suffice for our purposes; a considerably richer family is needed (see Remark 6.3).

et al. and the present paper are complementary and not directly comparable. We do note that (i) their guarantees are in expectation rather than with high-confidence, (ii) our TV-based metric is not covered by their smooth f-divergence family, and most important (iii) their notion of mixing is related to contraction as opposed to the spectral gap. In particular the α-minoration assumption implies (but is not implied by) a bound of $\kappa \leq 1 - d\alpha$ on the Dobrushin contraction coefficient (defined in (2.8); see Kontorovich [17], Lemma 2.2.2, for the latter claim). Thus, the family of α-minorized Markov chains is strictly contained in the family of contracting chains, which in turn is a strict subset of the ergodic chains we consider.

This paper is based on the conference version of Wolfer and Kontorovich [30] together with an extension to countably infinite spaces at Section 3.2. Another key improvement over the extended abstract is in the proof of Theorem 3.2. While the series of lemmas Wolfer and Kontorovich [30], Lemma 8, Lemma 9, Lemma 11, showed that it is possible to control the pseudo-spectral gap of our special family of chains via Cheeger's inequality combined with a contraction-based argument, this technique relied on heavy computations to bound the Dobrushin coefficient of the two-step transition matrix. Moreover, the proof for the extension to all members of the class was only sketched in Wolfer and Kontorovich [30], Lemma 9. In the present manuscript, we switch technique, compute the full spectrum of the unperturbed transition matrix instead, and fully flesh out the proof for the extension to perturbed chains using comparison techniques. Finally, the upper bound at Theorem 3.1 also improves upon Wolfer and Kontorovich [30], Theorem 1, by relying on a simulation scheme from Billingsley, instead of martingale techniques.

6. Proofs

6.1. Proof of Theorem 3.1

Remark 6.1. We thank an anonymous referee for the suggestion (and technique) to improve the logarithmic gap between the upper and lower bounds.

Let $\varepsilon \in (0, 2)$, $\delta \in (0, 1)$, let M be a d-state ergodic Markov kernel with stationary distribution π, and first consider the stationary case $X_1, \ldots, X_m \sim (M, \pi)$. We define the natural counting random variables

$$N_i \doteq \sum_{t=1}^{m-1} \mathbf{1}\{X_t = i\}, \qquad N_{ij} \doteq \sum_{t=1}^{m-1} \mathbf{1}\{X_t = i, X_{t+1} = j\}$$

and the estimator of the kernel will be $\widehat{M}(i, j) \doteq \frac{N_{ij}}{N_i}$ when $N_i \neq 0$ and $1/d$ when $N_i = 0$. We decompose the error probability of the estimation procedure, while choosing an arbitrary value $n_i \in \mathbb{N}$ for the desired number of visits to each state $i \in [d]$,

$$\mathbb{P}_{\pi}\left(\|M - \widehat{M}\|_{\infty} > \varepsilon\right) \leq \sum_{i=1}^{d} \sum_{n=n_i}^{3n_i} \mathbb{P}_{\pi}\left(\|\widehat{M}(i, \cdot) - M(i, \cdot)\|_1 > \varepsilon \text{ and } N_i = n\right)$$

$$+ \mathbb{P}_{\pi}\left(\{\exists i \in [d] : N_i \notin [n_i, 3n_i]\}\right). \tag{6.1}$$

We simulate a trajectory from M with a collection of independent samples using the scheme described in (Billingsley [4], p.19), where we define the following infinite array of random variables,

$$
\begin{array}{ccccc}
X_{1,1} & X_{1,2} & \cdots & X_{1,t} & \cdots \\
X_{2,1} & X_{2,2} & \cdots & X_{2,t} & \cdots \\
\cdots & \cdots & \cdots & \cdots & \cdots \\
X_{d,1} & X_{d,2} & \cdots & X_{d,t} & \cdots
\end{array}
$$

such that $\forall (i, j, t) \in [d]^2 \times \mathbb{N}$, $\mathbb{P}(X_{i,t} = j) = M(i, j)$, and the sampling procedure is as follows. Start by drawing $\tilde{X}_1 \sim \pi$. $\tilde{X}_2$ is then defined to be $X_{\tilde{X}_1, 1}$, the first element of the $\tilde{X}_1$th row. The process then continues inductively recording random variables from left to right in their corresponding rows, such that if $\tilde{X}_1, \tilde{X}_2, \ldots, \tilde{X}_t$ have been defined, then $\tilde{X}_{t+1} \doteq X_{\tilde{X}_t, \tilde{N}_{\tilde{X}_t}^{(t)} + 1}$ where $\tilde{N}_i^{(t)} \doteq \sum_{s=1}^t \mathbf{1}\{\tilde{X}_s = i\}$, and for convenience, $\tilde{N}_i \doteq \tilde{N}_i^{(m-1)}$. Observe that $X_1, X_2, \ldots,$ and $\tilde{X}_1, \tilde{X}_2, \ldots$ are identically distributed. Then, writing

$$
\widetilde{M}(i, \cdot) \doteq \frac{1}{\tilde{N}_i} \sum_{j=1}^d \sum_{t=1}^{m-1} \mathbf{1}\{\tilde{X}_t = i, \tilde{X}_{t+1} = j\} e_j,
$$

we have

$$
\mathbb{P}_\pi \left(\left\| \widehat{M}(i, \cdot) - M(i, \cdot) \right\|_1 > \varepsilon \text{ and } N_i = n \right)
$$
$$
= \mathbb{P} \left(\left\| \widetilde{M}(i, \cdot) - M(i, \cdot) \right\|_1 > \varepsilon \text{ and } \tilde{N}_i = n \right).
$$

In the event where $\tilde{N}_i = n$,

$$
\widetilde{M}(i, \cdot) = \frac{1}{n} \sum_{j=1}^d \sum_{t=1}^{m-1} \mathbf{1}\{\tilde{X}_t = i, X_{\tilde{X}_t, \tilde{N}_{\tilde{X}_t}^{(t)} + 1} = j\} e_j
$$

$$
= \frac{1}{n} \sum_{j=1}^d \sum_{t=1}^n \mathbf{1}\{X_{i,t} = j\} e_j
$$

$$
\doteq \widetilde{M}_n(i, \cdot),
$$

where by definition, $X_{i,1}, X_{i,2}, \ldots, X_{i,n} \sim M(i, \cdot)^{\otimes n}$, and the problem is reduced to learning a distribution out of n independent samples. Since $\mathbb{E}\|\widetilde{M}_n(i, \cdot) - M(i, \cdot)\|_1 \leq \sqrt{\frac{d}{n}}$, (see, for example, Berend and Kontorovich [3]) and the function $X_{i,1}, X_{i,2}, \ldots, X_{i,n} \mapsto \|\widetilde{M}_n(i, \cdot) - M(i, \cdot)\|_1$ is $(2/n)$-Lipschitz, an application of McDiarmid's inequality yields that

$$
\mathbb{P}\left(\left\| \widetilde{M}(i, \cdot) - M(i, \cdot) \right\|_1 > \varepsilon \text{ and } \tilde{N}_i = n \right)
$$

$$
\leq \mathbb{P}\left(\left\| \widetilde{M}_n(i, \cdot) - M(i, \cdot) \right\|_1 > \varepsilon \right)
$$

$$
\leq \exp\left(-\frac{n}{2} \max\left\{ 0, \varepsilon - \sqrt{\frac{d}{n}} \right\}^2 \right).
$$

It follows that

$$\sum_{n=n_i}^{3n_i} \mathbb{P}_{\boldsymbol{\pi}} \left(\left\| \widehat{\boldsymbol{M}}(i, \cdot) - \boldsymbol{M}(i, \cdot) \right\|_1 > \varepsilon \text{ and } N_i = n \right)$$

$$\overset{(i)}{\leq} (2n_i + 1) \exp\left(-\frac{n_i}{2} \max\left\{ 0, \varepsilon - \sqrt{\frac{d}{n_i}} \right\}^2 \right)$$

$$\overset{(ii)}{\leq} (m\pi_i + 1) \exp\left(-cm\pi_i \varepsilon^2 \right),$$

where $c = \frac{(1-1/\sqrt{2})^2}{4}$, (i) stems from a monotonicity argument, and (ii) is by setting $n_i = \frac{m\pi_i}{2}$, and as long as $m \geq \frac{4d}{\varepsilon^2 \pi_\star}$. We start by handling the first term of (6.1),

$$\sum_{i=1}^{d}(m\pi_i + 1) \exp\left(-cm\pi_i \varepsilon^2 \right) \leq \sum_{i=1}^{d} \frac{1}{c\varepsilon^2} \exp\left(-cm\pi_i \varepsilon^2/2 \right) + \sum_{i=1}^{d} \exp\left(-cm\pi_i \varepsilon^2 \right)$$

$$\leq \sum_{i=1}^{d} \frac{2}{c\varepsilon^2} \exp\left(-cm\pi_i \varepsilon^2/2 \right) \tag{6.2}$$

$$\leq \frac{2d}{c\varepsilon^2} \exp\left(-cm\pi_\star \varepsilon^2/2 \right),$$

where we used the fact that $x > 0 \implies x \exp(-x) \leq \exp(-x/2)$, and which is smaller than $\delta/2$ as long as $m \geq \frac{2}{c\varepsilon^2 \pi_\star} \ln \frac{4d}{c\delta \varepsilon^2}$. It remains to control the probability of the bad event where the states are not visited a reasonable amount of time. Invoking Paulin [22], Theorem 3.10,

$$\mathbb{P}_{\boldsymbol{\pi}} \left(N_i \notin \left[\frac{1}{2}m\pi_i, \frac{3}{2}m\pi_i \right] \right) \leq \exp\left(-\frac{\gamma_{\mathsf{ps}}(\frac{1}{2}m\pi_i)^2}{8(m + 1/\gamma_{\mathsf{ps}})\pi_i(1 - \pi_i) + 20\frac{1}{2}m\pi_i} \right). \tag{6.3}$$

Quantifying the price for non-stationarity using Paulin [22], Proposition 3.14,

$$\mathbb{P}_{\boldsymbol{\mu}} \left(\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\|_\infty > \varepsilon \right) \leq \|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}} \sqrt{\mathbb{P}_{\boldsymbol{\pi}} \left(\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\|_\infty > \varepsilon \right)},$$

and combining with (6.3) yields the upper bound. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Remark 6.2. Note that one can derive an upper bound of $\frac{1}{\pi_\star} \max\{1/\varepsilon^2, 1/\gamma_{\mathsf{ps}}\}$ (up to logarithmic factors) for the problem with respect to the max norm $\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_{\mathsf{MAX}} = \max_{(i,j)\in[d]^2} |\boldsymbol{M}(i, j) - \widehat{\boldsymbol{M}}(i, j)|$. Similarly, for $p \in [1, 2)$, we can derive the more general upper bound (up to logarithmic factors)

$$\frac{1}{\pi_\star} \max\left\{ \frac{d^{2/p-1}}{\varepsilon^2}, \frac{1}{\gamma_{\mathsf{ps}}} \right\}$$

for the problem with respect to the norm $\|\boldsymbol{M} - \widehat{\boldsymbol{M}}\|_{\infty,p} \doteq \max_{i\in[d]} \|\boldsymbol{M}(i, \cdot) - \widehat{\boldsymbol{M}}(i, \cdot)\|_p$.

6.2. Proof of Theorem 3.3

Sample $X_1, \ldots, X_m \sim (M, \mu)$ where M is ergodic with stationary distribution π. We define the estimator $\widehat{Q} \doteq \frac{1}{m-1} N_{ij}$ with $N_{ij} \doteq \sum_{t=1}^{m-1} \mathbf{1}\{X_t = i, X_{t+1} = j\}$. We first focus on the stationary case where $\mu = \pi$.

Bounding the distance in expectation

From Jensen's inequality and stationarity,

$$\mathbb{E}_\pi \|\widehat{Q} - Q\| = \sum_{(i,j) \in \Omega^2} \mathbb{E}_\pi \left| \frac{N_{ij}}{m-1} - \pi(i) M(i, j) \right|$$

$$\leq \frac{1}{m-1} \sum_{(i,j) \in \Omega^2} \sqrt{\mathbb{E}_\pi \left[(N_{ij} - (m-1)\pi(i) M(i, j))^2 \right]}$$

$$= \frac{1}{m-1} \sum_{(i,j) \in \Omega^2} \sqrt{\mathbb{Var}_\pi [N_{ij}]},$$

and we are left with controlling a variance term. The next lemma defines a new Markov chain from $X_1, \ldots, X_m$ with an approximately similar mixing time.

Lemma 6.1. *Let $X = (X_1, \ldots, X_m) \sim (M, \pi)$ with mixing time t_{mix} and stationary distribution π, then $Y = ((X_1, X_2), (X_2, X_3), \ldots, (X_{m-1}, X_m))$ is also a finite state Markov chain with mixing time at most $t_{\mathrm{mix}} + 1$, with kernel $\widetilde{M} : \Omega^2 \times \Omega^2 \to [0, 1]$ and stationary distribution $\tilde{\pi}$, such that for all $(i, j, k, \ell) \in \Omega^4$,*

$$\widetilde{M}((i, j), (k, \ell)) = \mathbf{1}\{k = j\} M(k, \ell)$$

$$\tilde{\pi}((i, j)) = Q(i, j).$$

Proof. Let $X = (X_1, \ldots, X_m) \sim (M, \pi)$ with mixing time t_{mix} and stationary distribution π, we first show that $Y = ((X_1, X_2), (X_2, X_3), \ldots, (X_{m-1}, X_m))$ is also a finite state Markov chain. For all $t \in \mathbb{N}$ and $y = (y_1, \ldots, y_{t-1})$ with $y_s = (x_s, x_{s+1})$ for $s \in [t-2]$, and whenever defined,

$$\mathbb{P}\left(Y_t = (i, j) \mid Y_1 = y_1, \ldots, Y_{t-1} = y_{t-1} \right)$$

$$= \mathbb{P}\left((X_t, X_{t+1}) = (i, j) \mid X_1 = x_1, \ldots, X_{t-1} = x_{t-1}, X_t = x_t \right)$$

$$= \mathbf{1}\{i = x_t\} \mathbb{P}(X_{t+1} = j \mid X_t = x_t) = \mathbf{1}\{i = x_t\} M(i, j)$$

$$= \mathbb{P}\left(Y_t = (i, j) \mid Y_{t-1} = y_{t-1} \right),$$

which confirms the Markov property. Additionally, setting $\tilde{\pi}((i, j)) = Q(i, j)$,

$$\sum_{(i,j) \in \Omega^2} \tilde{\pi}((i, j)) \widetilde{M}((i, j), (k, \ell)) = \sum_{(i,j) \in \Omega^2} \tilde{\pi}((i, j)) \mathbf{1}\{k = j\} M(k, \ell) = M(k, \ell) \sum_{i \in \Omega} \tilde{\pi}((i, k))$$

$$= M(k, \ell) \sum_{i \in \Omega} \pi(i) M(i, k) = M(k, \ell) \pi(k) = \tilde{\pi}((k, \ell)).$$

This entails that Y is a Markov chain $(\tilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{M}})$ over the state space $\Omega \times \Omega$, and stationary distribution $\tilde{\boldsymbol{\pi}}$. Let $t \geq t_{\text{mix}} + 1$, and $\tilde{\delta}_{(i_1, j_1)}$ the distribution on $\Omega \times \Omega$ that puts mass 1 at (i_1, j_1) and 0 everywhere else, then

$$\big\| \tilde{\delta}_{(i_1, j_1)} \widetilde{\boldsymbol{M}}^t - \tilde{\boldsymbol{\pi}} \big\|_{\text{TV}} = \sum_{(i,j) \in \Omega^2} \big| \tilde{\delta}_{(i_1, j_1)} \widetilde{\boldsymbol{M}}^t (i, j) - \tilde{\boldsymbol{\pi}}(i, j) \big|.$$

One one hand,

$$\begin{aligned}
\tilde{\delta}_{(i_1, j_1)} \widetilde{\boldsymbol{M}}^t (i, j) &= \mathbb{P}\big(Y_{t+1} = (i, j) \mid Y_1 = (i_1, j_1)\big) \\
&= \mathbb{P}\big((X_{t+1}, X_{t+2}) = (i, j) \mid (X_1, X_2) = (i_1, j_1)\big) \\
&= \mathbb{P}(X_{t+2} = j \mid X_{t+1} = i)\mathbb{P}(X_{t+1} = i \mid X_2 = j_1) \\
&= \boldsymbol{M}(i, j)\boldsymbol{M}^{t-1}(j_1, i),
\end{aligned}$$

so that

$$\begin{aligned}
\big\| \tilde{\delta}_{(i_1, j_1)} \widetilde{\boldsymbol{M}}^t - \tilde{\boldsymbol{\pi}} \big\|_{\text{TV}} &= \frac{1}{2} \sum_{(i,j) \in \Omega^2} \big| \boldsymbol{M}(i, j)\boldsymbol{M}^{t-1}(j_1, i) - \boldsymbol{M}(i, j)\boldsymbol{\pi}(i) \big| \\
&= \frac{1}{2} \sum_{i \in \Omega} \big| \boldsymbol{M}^{t-1}(j_1, i) - \boldsymbol{\pi}(i) \big| = \big\| \delta_{j_1} \boldsymbol{M}^{t-1} - \boldsymbol{\pi} \big\|_{\text{TV}} \leq 1/4,
\end{aligned}$$

by definition of t_{mix}, and by the condition on t. $\qquad \square$

Corollary 6.1. *Let $X_1, \dots, X_m \sim (\boldsymbol{M}, \boldsymbol{\pi})$ and define the chain $Y \sim (\widetilde{\boldsymbol{M}}, \tilde{\boldsymbol{\pi}})$ from Lemma* 6.1, *then by Paulin* [22], *Theorem* 3.2, *Proposition* 3.4, *for $\phi : \Omega^2 \to \mathbb{R}^+$,*

$$\mathbb{V}\mathbf{ar}_{\tilde{\boldsymbol{\pi}}} \left[\sum_{t=1}^{m-1} \phi(Y_t) \right] \leq \frac{4m}{\gamma_{\text{ps}}} \mathbb{V}\mathbf{ar}_{\tilde{\boldsymbol{\pi}}}[\phi] \leq 8mt_{\text{mix}} \mathbb{V}\mathbf{ar}_{\tilde{\boldsymbol{\pi}}}[\phi],$$

and from the fact that

$$\mathbb{V}\mathbf{ar}_{\boldsymbol{\pi}} \big[\mathbf{1}\{X_t = i, X_{t+1} = j\} \big] = \boldsymbol{\pi}(i)\boldsymbol{M}(i, j)\big(1 - \boldsymbol{\pi}(i)\boldsymbol{M}(i, j)\big),$$

we get the following control on the variance term:

$$\mathbb{V}\mathbf{ar}_{\boldsymbol{\pi}}[N_{ij}] \leq 8mt_{\text{mix}} \boldsymbol{\pi}(i)\boldsymbol{M}(i, j)\big(1 - \boldsymbol{\pi}(i)\boldsymbol{M}(i, j)\big).$$

As a consequence of Corollary 6.1,

$$\mathbb{E}_{\boldsymbol{\pi}} \| \widehat{\boldsymbol{Q}} - \boldsymbol{Q} \| \leq 2 \sqrt{\frac{t_{\text{mix}} + 1}{m - 1} \| \boldsymbol{Q} \|_{1/2}}, \quad \text{where } \| \boldsymbol{Q} \|_{1/2} \doteq \left(\sum_{(i,j) \in \Omega^2} \sqrt{\boldsymbol{Q}(i, j)} \right)^2$$

and

$$m \geq 64 \frac{\| \boldsymbol{Q} \|_{1/2} t_{\text{mix}}}{\varepsilon^2} \quad \implies \quad \mathbb{E}_{\boldsymbol{\pi}} \| \widehat{\boldsymbol{Q}} - \boldsymbol{Q} \| \leq \varepsilon/2.$$

Bounding the fluctuations around the expectation

The strategy is to show that the loss is $(4/(m-1))$-Lipschitz, and simply invoke McDiarmid's inequality for Markov chains (Paulin [22], Corollary 2.10).

For $\boldsymbol{x} = (x_1, \ldots, x_m) \in \mathbb{R}^m$, let $\boldsymbol{x}^{(k)} = (x_1, \ldots, x_{k-1}, x_k', x_{k+1}, \ldots, x_m)$, with $k \in [m]$,

$$
\left| \||\widehat{\boldsymbol{Q}}(\boldsymbol{x}) - \boldsymbol{Q}\|| - \||\widehat{\boldsymbol{Q}}(\boldsymbol{x}') - \boldsymbol{Q}\|| \right|
$$

$$
= \frac{1}{m-1} \left| \sum_{(i,j) \in \Omega^2} \left(\left| n_{ij} - (m-1)\boldsymbol{Q}(i,j) \right| - \left| n_{ij}^{(k)} - (m-1)\boldsymbol{Q}(i,j) \right| \right) \right|
$$

$$
\leq \frac{1}{m-1} \sum_{(i,j) \in \Omega^2} \left| |n_{ij} - (m-1)\boldsymbol{Q}(i,j)| - |n_{ij}^{(k)} - (m-1)\boldsymbol{Q}(i,j)| \right|
$$

$$
\leq \frac{1}{m-1} \sum_{(i,j) \in \Omega^2} \left| n_{ij} - n_{ij}^{(k)} \right|,
$$

where we successively invoked the forward and reverse triangle inequality. We then compute

$$
\sum_{(i,j) \in \Omega^2} \left| n_{ij} - n_{ij}^{(k)} \right| = \sum_{(i,j) \in \Omega^2} \left| \mathbf{1}\{x_k = i, x_{k+1} = j\} - \mathbf{1}\{x_k^{(k)} = i, x_{k+1}^{(k)} = j\} \right.
$$

$$
\left. + \mathbf{1}\{x_{k-1} = i, x_k = j\} - \mathbf{1}\{x_{k-1}^{(k)} = i, x_k^{(k)} = j\} \right|
$$

$$
\leq \sum_{(i,j) \in \Omega^2} \left(\mathbf{1}\{x_k = i, x_{k+1} = j\} + \mathbf{1}\{x_k' = i, x_{k+1} = j\} \right.
$$

$$
\left. + \mathbf{1}\{x_{k-1} = i, x_k = j\} + \mathbf{1}\{x_{k-1} = i, x_k' = j\} \right) \leq 4,
$$

so that $\||\widehat{\boldsymbol{Q}}(\boldsymbol{x}) - \boldsymbol{Q}\|| - \||\widehat{\boldsymbol{Q}}(\boldsymbol{x}') - \boldsymbol{Q}\|| \leq \frac{4}{m-1}$, and it is then a consequence of McDiarmid's inequality for Markov chains (Paulin [22], Corollary 2.10) that

$$
\mathbb{P}_{\pi} \left(\left| \||\widehat{\boldsymbol{Q}}(\boldsymbol{X}) - \boldsymbol{Q}\|| - \mathbb{E}_{\pi} \||\widehat{\boldsymbol{Q}}(\boldsymbol{X}) - \boldsymbol{Q}\|| \right| > \varepsilon/2 \right) \leq 2 \exp\left(-\frac{m\varepsilon^2}{c t_{\mathrm{mix}}} \right), \quad c \in \mathbb{R}_+
$$

and $m \geq \frac{c t_{\mathrm{mix}}}{\varepsilon^2} \ln(\frac{2}{\delta}) \implies \mathbb{P}_{\pi} (|\||\widehat{\boldsymbol{Q}} - \boldsymbol{Q}\|| - \pi \||\widehat{\boldsymbol{Q}}(\boldsymbol{x}) - \boldsymbol{Q}\||| > \varepsilon/2) \leq \delta$. Finally, we extend the study to non-stationary chains in a straightforward way as for the proof of Theorem 6.1, with Paulin [22], Proposition 3.10, which yields the final theorem. $\qquad \square$

6.3. Proof of Theorem 3.2 (part 1): $\frac{d}{\varepsilon^2 \pi_\star}$

Recall the definition of KL divergence between two distributions $(\boldsymbol{v}, \boldsymbol{\theta}) \in \Delta_\Omega^2$, such that $\boldsymbol{v} \ll \boldsymbol{\theta}$,

$$
D_{\mathrm{KL}}(\boldsymbol{v} \| \boldsymbol{\theta}) \doteq \sum_{i \in \Omega} \boldsymbol{v}(i) \ln \frac{\boldsymbol{v}(i)}{\boldsymbol{\theta}(i)}.
$$

Let $\varepsilon \in (0, 1/32)$, and $\mathcal{M}_{d, \gamma_{\mathrm{ps}}, \pi_\star}$ be the collection of all d-state Markov chains whose stationary distribution is minorized by $\pi_\star$ and whose pseudo-spectral gap is at least γ_{ps}. The quantity we wish to lower

bound is the minimax risk for the estimation problem:

$$\mathcal{R}_m = \inf_{\widehat{M}} \sup_{M} \mathbb{P}_M \big(\|M - \widehat{M}\|_\infty > \varepsilon \big), \tag{6.4}$$

where the inf is taken over all estimation procedures $\widehat{M} : [d]^m \to \mathcal{M}_d$, $X \mapsto \widehat{M}(X)$ and the sup over $\mathcal{M}_{d,\gamma_{\mathrm{ps}},\pi_\star}$. Suppose for simplicity of the analysis that we consider Markov chains of $d + 1$ states instead of d, and that d is even. A slight modification of the proofs covers the odd case. We define the following class of Markov chains parametrized by a given distribution $p \in \Delta_{d+1}$, where the conditional distribution defined at each state of the chain is always p with $p_{d+1} = p_\star$ and $p_k = \frac{1-p_\star}{d}$ for $k \in [d]$, with $p_\star < \frac{1}{d+2}$, except for state $d + 1$, where it is only required that it has a loop of probability $p_\star$ to itself.

$$\mathcal{G}_p = \left\{ M_\eta = \begin{pmatrix} p_1 & \cdots & p_d & p_\star \\ \vdots & \vdots & \vdots & \vdots \\ p_1 & \cdots & p_d & p_\star \\ \eta_1 & \cdots & \eta_d & p_\star \end{pmatrix} : \eta = (\eta_1, \ldots, \eta_d, p_\star) \in \Delta_{d+1} \right\}. \tag{6.5}$$

Remark: a family of Markov chains very similar to $\mathcal{G}_p$ was independently considered by Hao et al. [10] for proving their lower bound.

It is easy to see that the stationary distribution π of an element of $\mathcal{G}_p$ indexed by η is

$$\pi_k = \frac{(1-p_\star)^2}{d} + \eta_k p_\star, \quad \text{for } k \in [d], \qquad \pi_{d+1} = p_\star.$$

For $m \geq 4$, $\eta = (\eta_1, \ldots, \eta_d, p_\star) \in \Delta_{d+1}$ and $(X_1, \ldots, X_m) \sim (M_\eta, p)$, set $N_i = \sum_{t=1}^m \mathbf{1}\{X_t = i\}$ the number of visits to the ith state. Focusing on the $(d + 1)$th state, since for $i \in [d + 1]$, we have $M_\eta(i, d + 1) = p_\star$, it is immediate that $N_{d+1} \sim \mathrm{Binomial}(m, p_\star)$. Introduce the subset of Markov chains in $\mathcal{G}_p$ such that

$$\eta(\sigma) = \left(\frac{1-p_\star + 16\sigma_1\varepsilon}{d}, \frac{1-p_\star - 16\sigma_1\varepsilon}{d}, \ldots, \frac{1-p_\star + 16\sigma_{\frac{d}{2}}\varepsilon}{d}, \frac{1-p_\star - 16\sigma_{\frac{d}{2}}\varepsilon}{d}, p_\star \right),$$

where $\sigma = (\sigma_1, \ldots, \sigma_{\frac{d}{2}}) \in \{-1, 1\}^{\frac{d}{2}}$. Also define M_0 with $\eta_0 = (\frac{1-p_\star}{d}, \ldots, \frac{1-p_\star}{d}, p_\star)$. We start by showing that for any chain of this family, γ_{ps} is bounded from below by a universal constant. The Dobrushin coefficient κ [defined at (2.8)] verifies

$$\kappa(M_\sigma) = \big\| \eta(\sigma) - \eta_0 \big\|_{\mathrm{TV}} = 8\varepsilon \leq 1/2.$$

From the Bubley-Dyer path coupling method (Bubley and Dyer [5]), $t_{\mathrm{mix}} \leq \frac{\ln 1/4}{\ln(1-\kappa)} \leq 2$, such that combining with Paulin [22], Proposition 3.4, $\gamma_{\mathrm{ps}} \geq \frac{1}{2t_{\mathrm{mix}}} \geq \frac{1}{4}$. A direct computation yields that for $\sigma \neq \sigma'$, $\|M_\sigma - M_{\sigma'}\|_1 = \frac{32\varepsilon}{d} d_H(\sigma, \sigma')$, where d_H is the Hamming distance. From the Varshamov-Gilbert lemma, we know that there is a $\Sigma \subset \{-1, 1\}^{d/2}$, $|\Sigma| \geq 2^{d/16}$, such that for $(\sigma, \sigma') \in \Sigma$ with $\sigma \neq \sigma'$, we have $d_H(\sigma, \sigma') \geq \frac{d}{16}$. Restricting our problem to this set Σ, and finally noticing that for $\sigma \in \Sigma$ we have $\|M_\sigma - M_0\|_1 = 16\varepsilon > 2\varepsilon$, applying Tsybakov's method (Tsybakov [25], Theorem 2.5) to our problem, we obtain

$$\mathcal{R}_m \geq \frac{1}{2} \left(1 - \frac{\frac{4}{2^{\frac{d}{16}}} \sum_{\sigma \in \Sigma} \mathcal{D}_m}{\ln 2^{\frac{d}{16}}} \right),$$

where $\mathcal{D}_m$ denotes the KL divergence between the two distributions of words of length m (see formal definition at Lemma 6.2) from each of the Markov chains indexed by $\eta(\sigma)$ and η_0. Leveraging the chain rule for the KL divergence, and as by construction, the only discrepancy occurs when visiting the $(d+1)$th state, Lemma 6.2 shows the following tensorization property,

$$\mathcal{D}_m \leq p_\star m D_{\mathrm{KL}}\big(\eta(\sigma)\|\eta_0\big), \tag{6.6}$$

following up with a straightforward computation,

$$D_{\mathrm{KL}}\big(\eta(\sigma)\|\eta_0\big) = \sum_{s\in\{-1,+1\}} \frac{d}{2}\left(\frac{1-p_\star+16s\varepsilon}{d}\right)\ln\left(\frac{\frac{1-p_\star+16s\varepsilon}{d}}{\frac{1-p_\star}{d}}\right) \leq 128\varepsilon^2, \tag{6.7}$$

and finally combining (6.4), (6.6) and (6.7), we get $\mathcal{R}_m \geq \frac{1}{2}(1 - \frac{512\varepsilon^2 mp_\star}{\frac{d}{16}\ln 2})$. Further noticing that for the considered range of ε and for $p_\star < \frac{1}{d+2}$, it is always the case that $\pi_\star = p_\star$, so that for $m \leq \frac{d(1-2\delta)\ln 2}{8192\varepsilon^2\pi_\star}$, $\mathcal{R}_m \geq \delta$.

Lemma 6.2. *For two Markov chains M_1 and M_2 of the class $\mathcal{G}_p$ defined at (6.5) indexed respectively by η_1 and η_2, denote respectively by $\mathcal{L}_1(X_1^m)$ and $\mathcal{L}_2(X_2^m)$ the distributions of words of length m, and write for simplicity $\mathcal{D}_t = D_{\mathrm{KL}}(\mathcal{L}_1(X_1^t)\|\mathcal{L}_2(X_2^t))$ the KL divergence between the processes up to time t. Then it is a fact that*

$$\mathcal{D}_m = (m-1)p_\star D_{\mathrm{KL}}(\eta_1\|\eta_2).$$

Proof. *From an application of the chain rule for the KL divergence, followed by the Markov property,*

$$\mathcal{D}_m = \mathcal{D}_{m-1} + \mathcal{E}_m,$$

where

$$\mathcal{E}_m \doteq \mathbb{E}_{X_1^{m-1}\sim\mathcal{L}_1}[\mathcal{D}_{m|m-1}],$$

$$\mathcal{D}_{m|m-1} \doteq D_{\mathrm{KL}}\big(\mathcal{L}_1(X_m \mid X_{m-1})\|\mathcal{L}_2(X_m \mid X_{m-1})\big).$$

In the event where $X_{m-1} \neq d+1$, $\mathcal{D}_{m|m-1} = 0$, such that from the law of total expectation,

$$\mathcal{E}_m = \mathbb{E}_{X_1^{m-2}\sim\mathcal{L}_1}\big[\mathbb{E}[\mathcal{D}_{m|m-1} \mid X_{m-1} = d+1]\mathcal{L}_1\big(X_{m-1} = d+1 \mid X_1^{m-2}\big)\big]. \tag{6.8}$$

From a second application of the Markov property, and by structural property of the chain,

$$\mathcal{L}_1\big(X_{m-1} = d+1 \mid X_1^{m-2}\big) = \mathcal{L}_1(X_{m-1} = d+1 \mid X_{m-2}) = M_1(X_{m-2}, d+1) = p_\star, \tag{6.9}$$

while in the event where $X_{m-1} = d+1$,

$$\mathcal{D}_{m|m-1} = D_{\mathrm{KL}}(\eta_1\|\eta_2). \tag{6.10}$$

Combining (6.8) with (6.9) and (6.10) yields,

$$\mathcal{E}_m = p_\star D_{\mathrm{KL}}(\eta_1\|\eta_2).$$

From an inductive argument, and the base case $\mathcal{D}_1 = 0$,

$$\mathcal{D}_m = (m-1)p_\star D_{\mathrm{KL}}(\eta_1\|\eta_2). \qquad \square$$

6.4. Proof of Theorem 3.2 (part 2): $\frac{d \ln d}{\gamma_{\mathsf{ps}}}$

We treat $\varepsilon \in (0, 1/8)$ and $d = 6k$, $k \geq 2$ as fixed. For $\eta \in (0, 1/48)$ and $\boldsymbol{\tau} \in \{0, 1\}^{d/3}$, define the block matrix

$$\boldsymbol{M}_{\eta, \boldsymbol{\tau}} = \begin{pmatrix} \boldsymbol{C}_\eta & \boldsymbol{R}_{\boldsymbol{\tau}} \\ \boldsymbol{R}_{\boldsymbol{\tau}}^{\mathsf{T}} & \boldsymbol{L}_{\boldsymbol{\tau}} \end{pmatrix},$$

where $\boldsymbol{C}_\eta \in \mathbb{R}^{d/3 \times d/3}$, $\boldsymbol{L}_{\boldsymbol{\tau}} \in \mathbb{R}^{2d/3 \times 2d/3}$, and $\boldsymbol{R}_{\boldsymbol{\tau}} \in \mathbb{R}^{d/3 \times 2d/3}$ are given by

$$\boldsymbol{L}_{\boldsymbol{\tau}} = \frac{1}{8} \operatorname{diag}(7 - 4\tau_1 \varepsilon, 7 + 4\tau_1 \varepsilon, \ldots, 7 - 4\tau_{d/3} \varepsilon, 7 + 4\tau_{d/3} \varepsilon),$$

$$\boldsymbol{C}_\eta = \begin{pmatrix} \frac{3}{4} - \eta & \frac{\eta}{d/3 - 1} & \cdots & \frac{\eta}{d/3 - 1} \\ \frac{\eta}{d/3 - 1} & \frac{3}{4} - \eta & \ddots & \vdots \\ \vdots & \ddots & \ddots & \frac{\eta}{d/3 - 1} \\ \frac{\eta}{d/3 - 1} & \cdots & \frac{\eta}{d/3 - 1} & \frac{3}{4} - \eta \end{pmatrix},$$

$$\boldsymbol{R}_{\boldsymbol{\tau}} = \frac{1}{8} \begin{pmatrix} 1 + 4\tau_1 \varepsilon & 1 - 4\tau_1 \varepsilon & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & 0 & 1 + 4\tau_2 \varepsilon & 1 - 4\tau_2 \varepsilon & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & \cdots & \cdots & 0 & 1 + 4\tau_{d/3} \varepsilon & 1 - 4\tau_{d/3} \varepsilon \end{pmatrix}.$$

Holding η fixed, define the collection

$$\mathcal{H}_\eta = \left\{ \boldsymbol{M}_{\eta, \boldsymbol{\tau}} : \boldsymbol{\tau} \in \{0, 1\}^{d/3} \right\} \tag{6.11}$$

of Markov matrices. Denote by $\boldsymbol{M}_{\eta, \boldsymbol{0}} \in \mathcal{H}_\eta$ the element corresponding to $\boldsymbol{\tau} = \boldsymbol{0}$. Note that every $\boldsymbol{M} \in \mathcal{H}_\eta$ is ergodic and reversible, and its unique stationary distribution is uniform. A graphical illustration of this class of Markov chains is provided in Figure 1; in particular, every $\boldsymbol{M} \in \mathcal{H}_\eta$ consists of an "inner clique" (i.e., the states indexed by $\{1, \ldots, d/3\}$) and "outer rim" (i.e., the states indexed by $\{d/3 + 1, \ldots, d\}$).

Lemma 6.3 in the Appendix establishes a key property of the elements of $\mathcal{H}_\eta$: each $\boldsymbol{M}$ in this class satisfies
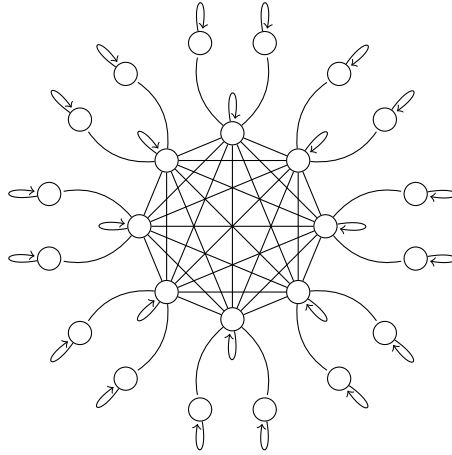
$$\eta/4 \leq \gamma_{\mathsf{ps}} \leq \eta.$$

Suppose that $X = (X_1, \ldots, X_m) \sim (\boldsymbol{M}_\eta, \boldsymbol{\pi})$, where $\boldsymbol{M} \in \mathcal{H}_\eta$ and $\boldsymbol{\pi}$ is uniform. Define the random variable T_{CLIQ}, to be the first time all of the states in the inner clique were visited,

$$T_{\mathrm{CLIQ}} = \inf\{t \geq 1 : \left|\{X_1, \ldots, X_t\} \cap [d/3]\right| = d/3\}, \tag{6.12}$$

Lemma 6.4 in the Appendix gives a lower estimate on this quantity:

$$m \leq \frac{d}{20\eta} \ln\left(\frac{d}{3}\right) \quad \implies \quad \mathbb{P}(T_{\mathrm{CLIQ}} > m) \geq \frac{1}{5}.$$

Figure 1. Generic topology of the $\mathcal{H}_\eta$ Markov chain class: every chain consists of an "inner clique" and an "outer rim".

Let $\mathcal{M}_{d,\gamma_{\mathrm{ps}},\pi_\star}$ be the collection of all d-state Markov chains whose stationary distribution is minorized by $\pi_\star$ and whose pseudo-spectral gap is at least γ_{ps}. Writing $X = (X_1, \ldots, X_m)$, recall that the quantity we wish to lower bound is the minimax risk for the statistical estimation problem (it will be convenient to write $\varepsilon/2$ instead of ε, which only affects the constants):

$$\mathcal{R}_m = \inf_{\widehat{M}} \sup_{M} \mathbb{P}_M\left(\|M - \widehat{M}\|_\infty > \frac{\varepsilon}{2} \right),$$

where the inf is taken over all estimation procedures $\widehat{M} : \Omega^m \to \mathcal{M}_d$, $X \mapsto \widehat{M}(X)$, and the sup over $\mathcal{M}_{d,\gamma_{\mathrm{ps}},\pi_\star}$. We employ the general reduction scheme of Tsybakov [25], Chapter 2.2. The first step is to restrict the sup to the finite subset $\mathcal{H}_\eta \subsetneq \mathcal{M}_{d,\gamma_{\mathrm{ps}},\pi_\star}$.

$$\mathcal{R}_m \geq \inf_{\widehat{M}} \max_{\tau} \mathbb{P}_{M_{\eta,\tau}}\left(\|M_{\eta,\tau} - \widehat{M}\|_\infty > \frac{\varepsilon}{2} \right).$$

Define T_{CLIQ} as in (6.12). Then

$$\mathcal{R}_m \geq \inf_{\widehat{M}} \max_{\tau} \mathbb{P}_{M_{\eta,\tau}}\left(\|M_{\eta,\tau} - \widehat{M}\|_\infty > \varepsilon \mid T_{\mathrm{CLIQ}} > m \right) \mathbb{P}_{M_{\eta,\tau}}(T_{\mathrm{CLIQ}} > m)$$

and Lemma 6.4 implies that for $m < \frac{d}{20\eta} \ln\left(\frac{d}{3}\right)$,

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\widehat{M}} \max_{\tau} \mathbb{P}_{M_{\eta,\tau}}\left(\|M_{\eta,\tau} - \widehat{M}\|_\infty > \varepsilon \mid T_{\mathrm{CLIQ}} > m \right).$$

Observe that all $\tau \neq \tau' \in \{0,1\}^{d/3}$ verify $\|M_{\eta,\tau} - M_{\eta,\tau'}\|_\infty = \varepsilon$. For any estimate $\widehat{M}$, define

$$\tau^\star = \operatorname*{argmin}_{\tau} \|\widehat{M} - M_{\eta,\tau}\|_\infty.$$

Then for $\tau \neq \tau^\star$, we have

$$\varepsilon = \|M_{\eta,\tau} - M_{\eta,\tau^\star}\|_\infty \leq \|M_{\eta,\tau} - \widehat{M}\|_\infty + \|\widehat{M} - M_{\eta,\tau^\star}\|_\infty \leq 2\|M_{\eta,\tau} - \widehat{M}\|_\infty,$$

whence $\{\tau^\star \neq \tau\} \subset \{\|M_{\eta,\tau} - \widehat{M}\|_\infty > \varepsilon/2\}$ and

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\widehat{M}} \max_\tau \mathbb{P}_{M_{\eta,\tau}} \left(\tau^\star \neq \tau \mid T_{\text{CLIQ}} > m \right)$$

$$= \frac{1}{5} \inf_{\hat{\tau}: X \mapsto \{0,1\}^{d/3}} \sup_\tau \mathbb{P}_{M_{\eta,\tau}} (\hat{\tau} \neq \tau \mid T_{\text{CLIQ}} > m).$$

Since $T_{\text{CLIQ}} > m$ implies that $N_i = 0$ for some $i \in [d/3]$,

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\hat{\tau}} \sup_\tau \mathbb{P}_{M_{\eta,\tau}} (\hat{\tau}_i \neq \tau_i \mid N_i = 0).$$

There are as many $M \in \mathcal{H}_\eta$ with $\tau_i = 0$ as those with $\tau_i = 1$, so if M is drawn uniformly at random and state i has not been visited, one can do no better than to make a random choice of $\hat{\tau}_i$ (where $\hat{\tau}$ determines $\widehat{M}$). More formally, writing $\tau^{(i)} = (\tau_1, \ldots, \tau_{i-1}, \tau_{i+1}, \ldots, \tau_{d/3}) \in \{0,1\}^{d/3-1}$, the τ vector without its ith coordinate, we can employ an Assouad-type of decomposition (Assouad [2], Yu [32]):

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\hat{\tau}} 2^{1-d/3} \sum_{\tau^{(i)} \in \{0,1\}^{d/3-1}} \left[\frac{1}{2} \mathbb{P}_{\tau_i=0} (\hat{\tau}_i \neq \tau_i \mid N_i = 0) + \frac{1}{2} \mathbb{P}_{\tau_i=1} (\hat{\tau}_i \neq \tau_i \mid N_i = 0) \right]$$

$$= \frac{2^{1-d/3}}{10} \sum_{\tau^{(i)} \in \{0,1\}^{d/3-1}} \inf_{\hat{\tau}} \left[\mathbb{P}_{\tau_i=0} (\hat{\tau}_i = 1 \mid N_i = 0) + \mathbb{P}_{\tau_i=1} (\hat{\tau}_i = 0 \mid N_i = 0) \right]$$

$$= \frac{2^{1-d/3}}{10} \sum_{\tau^{(i)} \in \{0,1\}^{d/3-1}} \left[1 - \left\| \mathbb{P}_{\tau_i=0} (X = \cdot \mid N_i = 0) + \mathbb{P}_{\tau_i=1} (X = \cdot \mid N_i = 0) \right\|_{\text{TV}} \right]$$

$$= \frac{1}{10}.$$

Combined with Lemma 6.3, and inclusion of events, this implies lower bound of $\frac{d}{\gamma_{\text{ps}}} \ln d$ for the estimation problem, which is tight for the case $\pi_\star = \frac{1}{d}$. □

Remark 6.3. Let us compare construction $\mathcal{H}_\eta$ to the family of Markov chains employed in the lower bound of Hsu et al. [12]:

$$M(i,j) = \begin{cases} 1 - \eta_i, & i = j, \\ \dfrac{\eta_i}{d-1}, & \text{else,} \end{cases}$$

where $\eta_i \in \{\eta, \eta'\}$ with $\eta' \approx \eta/2$. For our lower bound, $\mathcal{H}'_\eta$ has to be a ε-separated set under $\|\cdot\|_\infty$. In the construction of Hsu et al., the spectral gap γ and the separation distance ε are coupled, and using their family of Markov chains would lead to a lower bound of order $d/\gamma \approx d/\varepsilon$, which is inferior to $\frac{d}{\varepsilon^2 \pi_\star}$. The free parameter η was key to our construction, which enabled us to decouple γ from ε.

Lemma 6.3. *Let $\varepsilon \in (0, 1/32)$ and $\eta \in (0, 1/48)$. For all $M \in \mathcal{H}_\eta$ [defined in (6.11)], we have $\eta/4 \leq \gamma_\star \leq \gamma_{\text{ps}} \leq \eta$.*

Proof. We focus our proof on the absolute spectral gap, and will later show that the pseudo spectral gap is of the same order for our class of Markov matrices. A lower bound for $\gamma_\star$ of the unperturbed chain

$M_{\eta,\mathbf{0}}$, is given by Lemma 6.5. We now show how to extend to general $\boldsymbol{\tau}$ with comparison techniques. It is well known that (see, for example, Levin et al. [19], Lemma 13.7) that for a reversible chain M,

$$\gamma(M) = \min_{\substack{f:[d]\to\mathbb{R} \\ f\perp_{\boldsymbol{\pi}}\mathbf{1},\|f\|_2=1}} \mathcal{E}_M(f), \tag{6.13}$$

where

$$\mathcal{E}_M(f) \doteq \frac{1}{2}\sum_{(i,j)\in[d]^2} \left(f(i)-f(j)\right)^2\boldsymbol{\pi}(i)M(i,j)$$

is the Dirichlet form associated to M with stationary distribution $\boldsymbol{\pi}$. We now use this variational definition to control the spectral gap of the perturbed chains $M_{\eta,\boldsymbol{\tau}}$ in terms of the one of $M_{\eta,\mathbf{0}}$, relying on the fact that for both chains, the stationary distribution is uniform. Comparing transition matrices,

$$\begin{aligned}
\mathcal{E}_{M_{\eta,\boldsymbol{\tau}}}(f) &= \frac{1}{2}\sum_{(i,j)\in[d]^2}\left(f(i)-f(j)\right)^2\boldsymbol{\pi}(i)M_{\eta,\boldsymbol{\tau}}(i,j) \\
&\geq \frac{1}{2}\sum_{(i,j)\in[d]^2}\left(f(i)-f(j)\right)^2(1/d)(1-4\varepsilon)M_{\eta,\mathbf{0}}(i,j) \\
&= (1-4\varepsilon)\mathcal{E}_{M_{\eta,\mathbf{0}}}(f),
\end{aligned}$$

and by the definition at (6.13), $\gamma(M_{\eta,\boldsymbol{\tau}}) \geq (1-4\varepsilon)\gamma(M_{\eta,\mathbf{0}})$.

Extension to γ_{ps}

Now note that for a symmetric and lazy M, $\boldsymbol{\pi}$ is the uniform distribution, $M^\dagger = M^\mathsf{T} = M$, and $\gamma_{\mathsf{ps}} = \max_{k\geq 1}\{\frac{\gamma(M^{2k})}{k}\}$. Denoting by $1 = \lambda_1 > \lambda_2 \geq \cdots \geq \lambda_d$ the eigenvalues of M, we have that for all $i \in [d]$ and $k \geq 1$, λ_i^{2k} is an eigenvalue for M^{2k}, and furthermore $1 = \lambda_1^{2k} > \lambda_2^{2k} \geq \cdots \geq \lambda_d^{2k}$. Then

$$\gamma_{\mathsf{ps}} = \max_{k\geq 1}\frac{1-\lambda_2^{2k}}{k} = 1-\lambda_2^2$$

– that is, the maximum is achieved at $k = 1$. Indeed, $1-\lambda_2^{2k} = (1-\lambda_2^2)(\sum_{i=0}^{k-1}\lambda_2^{2i})$ and the latter sum is at most k since $\lambda_2^2 < \lambda_2 < 1$. As a result, $\gamma_{\mathsf{ps}} = 1-\lambda_2^2 = 1-(1-\gamma)^2 = \gamma(2-\gamma)$ and

$$\gamma \leq \gamma_{\mathsf{ps}} \leq 2\gamma,$$

which completes the proof. $\qquad\square$

Lemma 6.4 (Cover time). *For $M \in \mathcal{H}_\eta$ [defined in (6.11)], the random variable T_{CLIQ} [defined in (6.12)] satisfies*

$$m \leq \frac{d}{20\eta}\ln\left(\frac{d}{3}\right) \quad\Longrightarrow\quad \mathbb{P}(T_{\mathrm{CLIQ}} > m) \geq \frac{1}{5}$$

Proof. Let $M \in \mathcal{H}_\eta$ and $M_I \in \mathcal{M}_{d/3}$ be such that M_I consists only in the inner clique of M, and each outer rim state got absorbed into its unique inner clique neighbor:

$$M_I = \begin{pmatrix} 1 - \eta & \dfrac{\eta}{d/3 - 1} & \cdots & \dfrac{\eta}{d/3 - 1} \\ \dfrac{\eta}{d/3 - 1} & 1 - \eta & \ddots & \vdots \\ \vdots & \ddots & \ddots & \dfrac{\eta}{d/3 - 1} \\ \dfrac{\eta}{d/3 - 1} & \cdots & \dfrac{\eta}{d/3 - 1} & 1 - \eta \end{pmatrix}.$$

By construction, it is clear that T_{CLIQ} is almost surely greater than the cover time of M_I. The latter corresponds to a generalized coupon collection time $U = 1 + \sum_{i=1}^{d/3-1} U_i$ where U_i is the time increment between the ith and the $(i+1)$th unique visited state. Formally, if X is a random walk according to M_I (started from any state), then $U_1 = \min\{t > 1 : X_t \neq X_1\}$ and for $i > 1$,

$$U_i = \min\{t > 1 : X_t \notin \{X_1, \ldots, X_{U_{i-1}}\}\} - U_{i-1}.$$

The random variables $U_1, U_2, \ldots, U_{d/3-1}$ are independent and

$$U_i \sim \mathrm{Geometric}\left(\eta - \frac{(i-1)\eta}{d/3}\right),$$

whence

$$\mathbb{E}[U_i] = \frac{d/3}{\eta(d/3 - i + 1)}, \qquad \mathbb{V}\mathrm{ar}[U_i] = \left(1 - \left(\eta - \frac{(i-1)\eta}{d/3}\right)\right)\left(\eta - \frac{(i-1)\eta}{d/3}\right)^{-2}$$

and

$$\mathbb{E}[U] \geq 1 + \frac{d/3}{\eta}\sigma_{d/3-1}, \qquad \mathbb{V}\mathrm{ar}[U] \leq \frac{(d/3-1)^2}{\eta^2}\frac{\pi^2}{6}$$

where $\sigma_d = \sum_{i=1}^d \frac{1}{i}$, and $\pi = 3.1416\ldots$. Invoking the Paley–Zygmund inequality with $\theta = 1 - \frac{2\sqrt{2/3}}{\sigma_{d/3-1}}$ we have

$$\mathbb{P}\left(U > \theta \mathbb{E}[U]\right) \geq \left(1 + \frac{\mathbb{V}\mathrm{ar}[U]}{(1-\theta)^2(\mathbb{E}[U])^2}\right)^{-1} \geq \frac{1}{5}.$$

Further, $\sigma_{d/3-1} \geq \sigma_3 = 11/6$ implies

$$\theta \mathbb{E}[U] \geq \frac{3}{20} \cdot \frac{d/3}{\eta}\sigma_{d/3-1} \geq \frac{d}{20\eta}\ln\left(\frac{d}{3}\right),$$

and thus for $m \leq \frac{d}{20\eta}\ln\left(\frac{d}{3}\right)$, we have $\mathbb{P}(T_{\mathrm{CLIQ}} > m) \geq \frac{1}{5}$. $\square$

Lemma 6.5 (Spectrum of $M_{\eta,0} \in \mathcal{H}_\eta$). *Let $d = 6k$, $k \geq 2$ and $0 < \eta < 1/48$, and write $c_d \doteq d/(d-3)$. The spectrum of $M_{\eta,0}$ is*

$$\mathrm{Spec}(M_{\eta,0}) = \{\lambda_1, \lambda_+, \overline{\lambda}, \underline{\lambda}, \lambda_-\}$$

where $\lambda_1 = 1$ (*mult. 1*), $\lambda_{\pm} \doteq \frac{1}{16}(13 - 8\eta c_d \pm \sqrt{(3 + 8\eta c_d/3)^2 + 512\eta^2 c_d^2/9})$ (*each mult. $d/3 - 1$*) , $\underline{\lambda} = 5/8$ (*mult. 1*) , $\overline{\lambda} = 7/8$ (*mult. $d/3$*).

Moreover, $\lambda_{\star} \doteq \max_{\lambda \in \text{Spec}(M_{\eta,0})}\{|\lambda|, \lambda \neq 1\} = \lambda_{+}$, and $\eta/4 \leq \gamma_{\star} \leq \eta/2$.

Proof. By definition, and writing $c_d \doteq d/(d-3)$, $M_{\eta,0} = \begin{pmatrix} C_{\eta} & R_0 \\ R_0^{\mathsf{T}} & L_0 \end{pmatrix}$, where $L_0 = \frac{7}{8}I$,

$$C_{\eta} = \frac{\eta}{d/3-1}\mathbf{1}^{\mathsf{T}} \cdot \mathbf{1} - \left(\eta c_d - \frac{3}{4}\right)I \quad \text{and} \quad R_0 = \frac{1}{8}\begin{pmatrix} 1 & 1 & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & 0 & 1 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & \cdots & \cdots & 0 & 1 & 1 \end{pmatrix}.$$

As $M_{\eta,0}$ is a symmetric matrix, its spectrum is real. Let $\lambda \in \mathbb{R}$, and suppose first that $\lambda \neq 7/8$. In this case, $|L_0 - \lambda I| \neq 0$, and leveraging the block-structure of the matrix, it is a classical result (see for example Silvester [24]) that

$$|M_{\eta,0} - \lambda I| = |L_0 - \lambda I| \cdot |C_{\eta} - \lambda I - R_0(L_0 - \lambda I)^{-1}R_0^{\mathsf{T}}|.$$

A direct computation shows that

$$R_0(L_0 - \lambda I)^{-1}R_0^{\mathsf{T}} = \frac{1}{4(7 - 8\lambda)}I,$$

such that

$$C_{\eta} - \lambda I - R_0(L_0 - \lambda I)^{-1}R_0^{\mathsf{T}} = \frac{\eta}{d/3-1}\mathbf{1}^{\mathsf{T}} \cdot \mathbf{1} - \left(\eta c_d - \frac{3}{4} + \lambda + \frac{1}{4(7 - 8\lambda)}\right)I.$$

This implies that $|C_{\eta} - \lambda I - R_0(L_0 - \lambda I)^{-1}R_0^{\mathsf{T}}| = 0$ if and only if, $\eta c_d - \frac{3}{4} + \lambda + \frac{1}{4(7-8\lambda)} \in$ $\text{Spec}(\eta/(d/3 - 1)\mathbf{1}^{\mathsf{T}} \cdot \mathbf{1}) = \{0, \eta c_d\}$ where 0 has multiplicity $d/3 - 1$ and ηc_d has multiplicity 1. Let $\xi \geq 0$, then solutions for the equation $\xi + \lambda + \frac{1}{4(7-8\lambda)} = \frac{3}{4}$ are given by

$$\lambda_{\pm}(\xi) = \frac{13 - 8\xi \pm \sqrt{(3 + 8\xi/3)^2 + 512\xi^2/9}}{16}.$$

Setting $\xi = 0$ yields that $\lambda_1 = 1$ and $\underline{\lambda} = 5/8$ are eigenvalues both with multiplicity 1, while setting $\xi = \eta c_d$ yields that

$$\lambda_{\pm} = \frac{13 - 8\eta c_d \pm \sqrt{(3 + 8\eta c_d/3)^2 + 512\eta^2 c_d^2/9}}{16},$$

are both eigenvalues with multiplicity $d/3 - 1$. As the characteristic polynomial of $M_{\eta,0}$ has degree d, a natural consequence is that $\overline{\lambda} = 7/8$ is another eigenvalue with multiplicity $d/3$. It remains to order $\lambda_1, \lambda_-, \lambda_+, \underline{\lambda}, \overline{\lambda}$. Since $M_{\eta,0}$ is lazy, all eigenvalues are positive. Trivially, λ_1 is the largest eigenvalue, $\overline{\lambda} > \underline{\lambda}$ and always $\lambda_- \leq \lambda_+$. Additionally, $512\eta^2 c_d^2/9 \geq 0$ implies that $\lambda_+ \geq 1 - \eta c_d/3 \geq 1 - \eta/2$ for the considered range of d, which is in turn larger than $7/8$ for $\eta \leq 1/4$. As a result $\lambda_{\star} = \lambda_+$ and $\gamma_{\star} \leq \eta/2$. Furthermore, as one can write $\lambda_+ = \frac{1}{16}(13 - 8\eta c_d \pm \sqrt{(3 + 4\eta c_d/3)^2 - 8c_d\eta(1 - 8c_d\eta)})$, and since $1 - 8c_d\eta \geq 0$, $\lambda_+ \leq 1 - \eta/4$, whence $\gamma_{\star} \geq \eta/4$. $\square$

Acknowledgments

References

[1] Anthony, M. and Bartlett, P.L. (1999). *Neural Network Learning*: *Theoretical Foundations*. Cambridge: Cambridge Univ. Press. MR1741038 https://doi.org/10.1017/CBO9780511624216

[2] Assouad, P. (1983). Deux remarques sur l'estimation. *C. R. Acad. Sci. Paris Sér. I Math*. **296** 1021–1024. MR0777600

[3] Berend, D. and Kontorovich, A. (2013). A sharp estimate of the binomial mean absolute deviation with applications. *Statist. Probab. Lett*. **83** 1254–1259. MR3041401 https://doi.org/10.1016/j.spl.2013.01.023

[4] Billingsley, P. (1961). Statistical methods in Markov chains. *Ann. Math. Stat*. **32** 12–40. MR0123420 https://doi.org/10.1214/aoms/1177705136

[5] Bubley, R. and Dyer, M. (1997). Path coupling: A technique for proving rapid mixing in Markov chains. In *Proceedings* 38*th Annual Symposium on Foundations of Computer Science* 223–231. IEEE.

[6] Craig, B.A. and Sendi, P.P. (2002). Estimation of the transition matrix of a discrete-time Markov chain. *Health Econ*. **11** 33–42.

[7] Devroye, L. and Lugosi, G. (2001). *Combinatorial Methods in Density Estimation*. *Springer Series in Statistics*. New York: Springer. MR1843146 https://doi.org/10.1007/978-1-4613-0125-7

[8] Fill, J.A. (1991). Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *Ann. Appl. Probab*. **1** 62–87. MR1097464

[9] Han, Y., Jiao, J. and Weissman, T. (2015). Minimax estimation of discrete distributions under ℓ_1 loss. *IEEE Trans. Inf. Theory* **61** 6343–6354. MR3418968 https://doi.org/10.1109/TIT.2015.2478816

[10] Hao, Y., Orlitsky, A. and Pichapati, V. (2018). On learning Markov chains. In *Advances in Neural Information Processing Systems* 646–655.

[11] Horn, R.A. and Johnson, C.R. (1985). *Matrix Analysis*. Cambridge: Cambridge Univ. Press. MR0832183 https://doi.org/10.1017/CBO9780511810817

[12] Hsu, D., Kontorovich, A., Levin, D.A., Peres, Y., Szepesvári, C. and Wolfer, G. (2019). Mixing time estimation in reversible Markov chains from a single sample path. *Ann. Appl. Probab*. **29** 2439–2480. MR3983341 https://doi.org/10.1214/18-AAP1457

[13] Kamath, S., Orlitsky, A., Pichapati, D. and Suresh, A.T. (2015). On learning distributions from their samples. In *Proceedings of the* 28*th Conference on Learning Theory*, COLT 2015, *Paris*, *France*, *July* 3–6, 2015 1066–1100. Available at http://jmlr.org/proceedings/papers/v40/Kamath15.html.

[14] Kearns, M.J., Mansour, Y., Ron, D., Rubinfeld, R., Schapire, R.E. and Sellie, L. (1994). On the learnability of discrete distributions. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing* **1994** 273–282. Montréal, Québec, Canada. https://doi.org/10.1145/195058.195155

[15] Kemeny, J.G. and Snell, J.L. (1976). *Finite Markov Chains*. New York: Springer. Reprinting of the 1960 original, Undergraduate Texts in Mathematics. MR0410929

[16] Kontorovich, A. and Pinelis, I. (2019). Exact lower bounds for the agnostic probably-approximately-correct (PAC) machine learning model. *Ann. Statist*. **47** 2822–2854. MR3988774 https://doi.org/10.1214/18-AOS1766

[17] Kontorovich, A.L. (2007). Measure Concentration of Strongly Mixing Processes with Applications Ph.D. thesis, Carnegie Mellon University.

[18] Kontorovich, L. and Ramanan, K. (2008). Concentration inequalities for dependent random variables via the martingale method. *Ann. Probab*. **36** 2126–2158. MR2478678 https://doi.org/10.1214/07-AOP384

[19] Levin, D.A., Peres, Y. and Wilmer, E.L. (2009). *Markov Chains and Mixing Times*, 2nd ed. Providence: American Mathematical Soc.

[20] Markov, A.A. (1906). Extension of the law of large numbers to dependent quantities. *Izvestiia Fiz.-Matem. Obsch. Kazan Univ.* **15** 135–156.

[21] Orlitsky, A. and Suresh, A.T. (2015). Competitive distribution estimation: Why is good-Turing good. In *Advances in Neural Information Processing Systems* 28: *Annual Conference on Neural Information Processing Systems* 2143–2151. Available at http://papers.nips.cc/paper/5762-competitive-distribution-estimation-why-is-good-turing-good.

[22] Paulin, D. (2015). Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electron. J. Probab.* **20** 79. MR3383563 https://doi.org/10.1214/EJP.v20-4039

[23] Roberts, G.O. and Rosenthal, J.S. (1997). Geometric ergodicity and hybrid Markov chains. *Electron. Commun. Probab.* **2** 13–25. MR1448322 https://doi.org/10.1214/ECP.v2-981

[24] Silvester, J.R. (2000). Determinants of block matrices. *Math. Gaz.* **84** 460–467. Available at http://www.jstor.org/stable/3620776.

[25] Tsybakov, A.B. (2009). *Introduction to Nonparametric Estimation*. *Springer Series in Statistics*. New York: Springer. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats. MR2724359 https://doi.org/10.1007/b13794

[26] Vidyasagar, M. (2014). An elementary derivation of the large deviation rate function for finite state Markov chains. *Asian J. Control* **16** 1–19. MR3160167 https://doi.org/10.1002/asjc.806

[27] Waggoner, B. (2015). ℓ_p testing and learning of discrete distributions. In *ITCS'15 – Proceedings of the 6th Innovations in Theoretical Computer Science* 347–356. New York: ACM. MR3419028 https://doi.org/10.1145/2688073.2688095

[28] Welton, N. and Ades, A.E. (2005). Estimation of Markov chain transition probabilities and rates from fully and partially observed data: Uncertainty propagation, evidence synthesis, and model calibration. *Med. Decis. Mak.* **25** 633–645.

[29] Wolfer, G. (2020). Mixing time estimation in ergodic Markov chains from a single trajectory with contraction methods. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory* (A. Kontorovich and G. Neu, eds.). *Proceedings of Machine Learning Research* **117** 890–905. San Diego, CA, USA. Available at http://proceedings.mlr.press/v117/wolfer20a.html.

[30] Wolfer, G. and Kontorovich, A. (2019). Minimax learning of ergodic Markov chains. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, *Proceedings of Machine Learning Research* **98** 904–930. Chicago, IL. Available at http://proceedings.mlr.press/v98/wolfer19a.html.

[31] Wolfer, G. and Kontorovich, A. (2019). Estimating the mixing time of ergodic Markov chains. In *Proceedings of the Thirty-Second Conference on Learning Theory*. *Proceedings of Machine Learning Research* **99** 3120–3159. Phoenix, USA. Available at http://proceedings.mlr.press/v99/wolfer19a.html.

[32] Yu, B. (1997). Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam* 423–435. New York: Springer. MR1462963