
Minimax Testing of Identity to a Reference Ergodic Markov Chain

Geoffrey Wolfer

Department of Computer Science
Ben-Gurion University of the Negev

Aryeh Kontorovich

Department of Computer Science
Ben-Gurion University of the Negev

Abstract

We exhibit an efficient procedure for testing, based on a single long state sequence, whether an unknown Markov chain is identical to or ε -far from a given reference chain. We obtain nearly matching (up to logarithmic factors) upper and lower sample complexity bounds for our notion of distance, which is based on total variation. Perhaps surprisingly, we discover that the sample complexity depends solely on the properties of the known reference chain and does not involve the unknown chain at all, which is not even assumed to be ergodic.

1 Introduction

Distinguishing whether an unknown *discrete* distribution \mathbf{D} is identical to a reference one $\bar{\mathbf{D}}$ or is ε -far from it in total variation (TV) is a special case of statistical *property testing*. For the iid case, it is known that a sample of size $\Theta(\sqrt{d}/\varepsilon^2)$, where d is the support size, is both sufficient and necessary (Batu et al., 2001; Valiant and Valiant, 2017). This is in contradistinction to the corresponding learning problem, with the considerably higher sample complexity of $\Theta(d/\varepsilon^2)$ (see, e.g., Anthony and Bartlett (1999); Waggoner (2015); Kontorovich and Pinelis (2019)). The Markovian setting has so far received no attention in the property testing framework, with the notable exception of the recent work of Daskalakis et al. (2018) and Cherapanamjeri and Bartlett (2019), which we discuss in greater detail in Section 2. Daskalakis et al. “initiate[d] the study of Markov chain testing”, but imposed the stringent constraint of being symmetric on both the reference and unknown chains. In this paper, we only require ergod-

icity of the reference chain, and make no assumptions on the unknown one — from which the tester receives a single long trajectory of observations — other than it having d states.

Our contribution. We prove nearly matching upper and lower bounds on the sample complexity of the testing problem in terms of the accuracy ε and the number of states d , as well as parameters derived from the stationary distribution and mixing. We discover that for testing, only the reference chain affects the sample complexity, and no assumptions (including ergodicity) need be made on the unknown one. In particular, we exhibit an efficient testing procedure, which, given a Markovian sequence of length

$$m = \tilde{O} \left(\frac{1}{\pi_\star} \max \left\{ \frac{\sqrt{d}}{\varepsilon^2}, t_{\text{mix}} \right\} \right), \quad (1.1)$$

correctly identifies the unknown chain with high probability, where $\pi_\star \leq 1/d$ and t_{mix} are, respectively, the minimum stationary probability and mixing time of the *known* reference chain. We also derive an instance-specific version of the previous bound: the factor $\frac{\sqrt{d}}{\pi_\star}$ in (1.1) can be replaced with the potentially much smaller quantity $\max_{i \in [d]} \left\{ \pi(i)^{-1} \|\bar{\mathbf{M}}(i, \cdot)\|_{2/3} \right\}$, defined in (4.1). Additionally, we construct two separate worst-case lower bounds of $\Omega(dt_{\text{mix}})$ and $\Omega\left(\frac{\sqrt{d}}{\pi_\star \varepsilon^2}\right)$, exhibiting a regime for which our testing procedure is unimprovable.

2 Related work

We consider *distribution testing* in the *property testing* (within the more classical statistical *hypothesis testing*¹) framework — a research program initiated by Batu et al. (2000).

The special case of iid uniformity testing was addressed (for various metrics) by Goldreich and Ron (2011);

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

¹A recent result in this vein is Barsotti et al. (2016); see also references therein.

Paninski (2008). Extensions to iid identity testing for arbitrary finite distributions were then obtained (Goldreich, 2016; Diakonikolas et al., 2019+), including the instance-optimal tester of Valiant and Valiant (2017), who showed that \sqrt{d} may be replaced with $\|\bar{\mathbf{D}}\|_{2/3}$, the $(2/3)$ -pseudo-norm of the reference distribution.

The study of goodness-of-fit testing for Markov chains can be traced back to Bartlett (1951); Hoel (1954); Anderson and Goodman (1957); Billingsley (1961). However, their approach is of an asymptotic nature, and does not provide with a finite-sample convergence analysis. To our knowledge, Daskalakis et al. (2018) were the first to consider the testing problem for Markov chains from within the distribution testing framework (see the introduction section therein for additional references and motivation for the testing problem). Their model is based on the pseudo-distance $\text{Dist}_{\sqrt{\cdot}}(\cdot, \cdot)$ defined by Kazakos (1978) as

$$\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 1 - \rho\left([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}\right),$$

where $\left([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}\right)_{(i,j)} = \sqrt{\mathbf{M}(i,j)\mathbf{M}'(i,j)}$ is the term-wise geometric mean of the transition kernels and ρ is the largest eigenvalue in magnitude. This pseudo-distance has the property of vanishing on pairs of chains sharing an identical connected component. Daskalakis et al.'s sample complexity upper bound of $\tilde{\mathcal{O}}(\text{HitT}_{\bar{\mathbf{M}}} + d/\varepsilon)$ required knowledge of the hitting time of the reference chain,

$$\text{HitT}_{\bar{\mathbf{M}}} = \max_{r,s \in \Omega} \mathbf{E}_{\bar{\mathbf{M}}}[\min\{t \geq 1: X_t = r\} | X_1 = s],$$

while their lower bound $\Omega(d/\varepsilon)$ involves no quantities related to the mixing rate at all. The authors conjectured that for $\text{Dist}_{\sqrt{\cdot}}(\cdot)$, the correct sample complexity is $\Theta(d/\varepsilon)$ — i.e., independent of the mixing properties of the chain. This conjecture was recently partially proven by Cherapanamjeri and Bartlett (2019), who gave an upper bound of $\tilde{\mathcal{O}}(d/\varepsilon^4)$, without dependence on the hitting time.

The present paper compares favorably with Daskalakis et al. in that the latter requires both the reference and the unknown chains to be symmetric (and, a fortiori, reversible) as well as ergodic. We only require ergodicity of the reference chain and assume nothing about the unknown one.

Additionally, we obtain nearly sharp sample complexity bounds in terms of the reference chain's mixing properties. Finally, our metric $\|\cdot\|$ dominates the pseudo-metric $\text{Dist}_{\sqrt{\cdot}}(\cdot)$, and hence their identity testing problem is reducible to ours (see Lemma 8.1 in the Appendix), although the reduction does not preserve the convergence rate.

We note that the corresponding PAC-type learning problem for Markov chains was only recently introduced in Hao et al. (2018); Wolfer and Kontorovich (2019). The present paper uses the same notion of distance as the latter work, in which the minimax sample complexity for learning was shown to be of the order of $m = \tilde{\Theta}\left(\frac{1}{\pi_{\star}} \max\left\{\frac{d}{\varepsilon^2}, t_{\text{mix}}\right\}\right)$. Our present results confirm the intuition that identity testing is, statistically speaking, a significantly less demanding task than learning: the former exhibits a quadratic reduction in the bound's dependence on d over the latter.

3 Definitions and notation

We define $[d] \doteq \{1, \dots, d\}$, denote the simplex of all distributions over $[d]$ by Δ_d , and the collection of all $d \times d$ row-stochastic matrices by \mathcal{M}_d . For $\boldsymbol{\mu} \in \Delta_d$, we will write either $\boldsymbol{\mu}(i)$ or μ_i , as dictated by convenience. All vectors are rows unless indicated otherwise. For $n \in \mathbb{N}$, and any $\mathbf{D} \in \Delta_d$ we also consider its n -fold product $\mathbf{D}^{\otimes n}$, i.e. $(X_1, \dots, X_n) \sim \mathbf{D}^{\otimes n}$ is a shorthand for $(X_i)_{i \in [n]}$ being all mutually independent, and such that $\forall i \in [n], X_i \sim \mathbf{D}$. A Markov chain on d states being entirely specified by an initial distribution $\boldsymbol{\mu} \in \Delta_d$ and a row-stochastic transition matrix $\mathbf{M} \in \mathcal{M}_d$, we identify the chain with the pair $(\mathbf{M}, \boldsymbol{\mu})$. Namely, writing X_1^m for (X_1, \dots, X_m) , by $X_1^m \sim (\mathbf{M}, \boldsymbol{\mu})$, we mean that

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

We write $\mathbf{P}_{\mathbf{M}, \boldsymbol{\mu}}(\cdot)$ to denote probabilities over sequences induced by the Markov chain $(\mathbf{M}, \boldsymbol{\mu})$, and omit the subscript when it is clear from context. Taking the null hypothesis to be that $\mathbf{M} = \bar{\mathbf{M}}$ (i.e., the chain being tested is identical to the reference one), $\mathbf{P}_0(\cdot)$ will denote probability in the completeness case, and $\mathbf{P}_1(\cdot)$ in the soundness case. The Markov chain $(\mathbf{M}, \boldsymbol{\mu})$ is *stationary* if $\boldsymbol{\mu} = \boldsymbol{\pi}$ for $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{M}$, and *ergodic* if $\mathbf{M}^k > 0$ (entry-wise positive) for some $k \geq 1$. If \mathbf{M} is ergodic, it has a unique stationary distribution $\boldsymbol{\pi}$ and moreover the *minimum stationary probability* $\pi_{\star} > 0$, where

$$\pi_{\star} \doteq \min_{i \in [d]} \boldsymbol{\pi}(i). \quad (3.1)$$

Unless noted otherwise, $\boldsymbol{\pi}$ is assumed to be the stationary distribution of the Markov chain in context. The *mixing time* of a chain is defined as the number of steps necessary for its state distribution to be sufficiently close to the stationary one (traditionally taken to be within $1/4$):

$$t_{\text{mix}} \doteq \min \left\{ t \geq 1 : \max_{\boldsymbol{\mu} \in \Delta_d} \|\boldsymbol{\mu}\mathbf{M}^t - \boldsymbol{\pi}\|_{\text{TV}} \leq \frac{1}{4} \right\}. \quad (3.2)$$

We use the standard ℓ_1 norm $\|z\|_1 = \sum_{i \in [d]} |z_i|$, which, in the context of distributions (and up to a convention-dependent factor of 2) corresponds to the total variation norm. For $\mathbf{A} \in \mathbb{R}^{d \times d}$, define

$$\|\mathbf{A}\| \doteq 2 \max_{i \in [d]} \|\mathbf{A}(i, \cdot)\|_{\text{TV}}.$$

Finally, we use standard $\mathcal{O}(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ order-of-magnitude notation, as well as their tilde variants $\tilde{\mathcal{O}}(\cdot)$, $\tilde{\Omega}(\cdot)$, $\tilde{\Theta}(\cdot)$ where lower-order log factors in any parameter are suppressed.

Definition 3.1 An (ε, δ) -identity tester \mathcal{T} for Markov chains with sample complexity function $m_0(\cdot)$ is an algorithm that takes as input a reference Markov chain $(\bar{\mathbf{M}}, \bar{\boldsymbol{\mu}})$ and $\mathbf{X} = (X_1, \dots, X_m)$ drawn from some unknown Markov chain $(\mathbf{M}, \boldsymbol{\mu})$, and outputs $\mathcal{T} = \mathcal{T}(d, \varepsilon, \delta, \bar{\mathbf{M}}, \bar{\boldsymbol{\mu}}, \mathbf{X}) \in \{0, 1\}$ such that for $m \geq m_0(d, \varepsilon, \delta, \bar{\mathbf{M}}, \bar{\boldsymbol{\mu}})$, both $\mathbf{M} = \bar{\mathbf{M}} \Rightarrow \mathcal{T} = 0$ and $\|\mathbf{M} - \bar{\mathbf{M}}\| > \varepsilon \Rightarrow \mathcal{T} = 1$ hold with probability at least $1 - \delta$. (The probability is over the draw of \mathbf{X} and any internal randomness of the tester.)

4 Formal results

Since the focus of this paper is on statistical rather than computational complexity, we defer the (straightforward) analysis of the runtimes of our tester to the Appendix, Section 7.

Theorem 4.1 (Upper bound) *There exists an (ε, δ) -identity tester \mathcal{T} (provided at Algorithm 1), which, for all $0 < \varepsilon < 2$, $0 < \delta < 1$, satisfies the following. If \mathcal{T} receives as input a d -state “reference” ergodic Markov chain $(\bar{\mathbf{M}}, \bar{\boldsymbol{\mu}})$, as well as a sequence $\mathbf{X} = (X_1, \dots, X_m)$ of length at least m_{UB} , drawn according to an unknown chain \mathbf{M} (starting from an arbitrary state), then it outputs $\mathcal{T} = \mathcal{T}(d, \varepsilon, \delta, \bar{\mathbf{M}}, \bar{\boldsymbol{\mu}}, \mathbf{X}) \in \{0, 1\}$ such that*

$$\begin{aligned} \mathbf{M} = \bar{\mathbf{M}} &\implies \mathcal{T} = 0 \\ \|\mathbf{M} - \bar{\mathbf{M}}\| > \varepsilon &\implies \mathcal{T} = 1 \end{aligned}$$

holds with probability at least $1 - \delta$. The sample complexity is upper-bounded by

$$\begin{aligned} m_{\text{UB}} &= \frac{C_{\text{MC}}}{\pi_\star} \max \left\{ \frac{\sqrt{d}}{\varepsilon^2} \ln \left(\frac{d}{\delta \varepsilon} \right), t_{\text{mix}} \ln \left(\frac{d}{\delta \pi_\star} \right) \right\} \\ &= \tilde{\mathcal{O}} \left(\frac{1}{\pi_\star} \max \left\{ \frac{\sqrt{d}}{\varepsilon^2}, t_{\text{mix}} \right\} \right), \end{aligned}$$

where C_{MC} is a universal constant, and t_{mix} and π_\star are respectively the mixing time (3.2) and the minimum stationary probability (3.1) of $\bar{\mathbf{M}}$.

Remark 4.1 *An important feature of Theorem 4.1 is that the sample complexity only depends on the (efficiently computable, see Section 7) properties of the known reference chain. No assumptions, such as symmetry (as in Daskalakis et al. (2018); Cherapanamjeri and Bartlett (2019)) or even ergodicity, are made on the unknown Markov chain, and none of its properties appear in the bound.*

Remark 4.2 *Our results indicate that in the regime where the $\frac{t_{\text{mix}}}{\pi_\star}$ term is not dominant, the use of optimized identity iid testers as subroutines confers an $\tilde{\mathcal{O}}(\sqrt{d})$ -fold improvement over the naive testing-by-learning strategy.*

Theorem 4.2 (Instance-specific upper bound)

There exists an (ε, δ) -identity tester \mathcal{T} , which, for all $0 < \varepsilon < 2$, $0 < \delta < 1$, satisfies the following. If \mathcal{T} receives as input a d -state “reference” ergodic Markov chain $(\bar{\mathbf{M}}, \bar{\boldsymbol{\mu}})$, as well as a sequence $\mathbf{X} = (X_1, \dots, X_m)$ of length at least m_{UB} , drawn according to an unknown chain \mathbf{M} (starting from an arbitrary state), then it outputs $\mathcal{T} = \mathcal{T}(d, \varepsilon, \delta, \bar{\mathbf{M}}, \bar{\boldsymbol{\mu}}, \mathbf{X}) \in \{0, 1\}$ such that

$$\begin{aligned} \mathbf{M} = \bar{\mathbf{M}} &\implies \mathcal{T} = 0 \\ \|\mathbf{M} - \bar{\mathbf{M}}\| > \varepsilon &\implies \mathcal{T} = 1 \end{aligned}$$

holds with probability at least $1 - \delta$. The sample complexity is upper-bounded by

$$m_{\text{UB}} = \tilde{\mathcal{O}} \left(\max \left\{ \frac{\|\bar{\mathbf{M}}\|_{\pi, 2/3}}{\varepsilon^2}, \frac{t_{\text{mix}}}{\pi_\star} \right\} \right),$$

where t_{mix} and π_\star are as in Theorem 4.1, and

$$\|\bar{\mathbf{M}}\|_{\pi, 2/3} \doteq \max_{i \in [d]} \left\{ \frac{\left(\sum_{j \in [d]} \bar{\mathbf{M}}(i, j)^{2/3} \right)^{3/2}}{\pi(i)} \right\}. \quad (4.1)$$

Remark 4.3 *Since we always have $\|\bar{\mathbf{M}}\|_{\pi, 2/3} \leq \sqrt{d}/\pi_\star$, the instance-specific bound is always at least as sharp as the worst-case one in Theorem 4.1. It may, however, be considerably sharper, as illustrated by a simple random walk on a d -vertex, Δ -regular graph (Levin et al., 2009, Section 1.4), for which the instance-specific bound is $\tilde{\mathcal{O}} \left(d \max \left\{ \frac{\sqrt{\Delta}}{\varepsilon^2}, t_{\text{mix}} \right\} \right)$ — a savings of roughly \sqrt{d} .*

Theorem 4.3 (Lower bounds) *For every $0 < \varepsilon < 1/8$, $t_{\text{mix}} \geq 50$, and $d = 6k$, $k \geq 2$, there exists a d -state Markov chain $\bar{\mathbf{M}}$ with mixing time t_{mix} and stationary distribution $\boldsymbol{\pi}$ such that every $(\varepsilon, 1/10)$ -identity tester for reference chain $\bar{\mathbf{M}}$ must require in the worst case a sequence $\mathbf{X} = (X_1, \dots, X_m)$ drawn from the unknown*

chain \mathbf{M} of length at least

$$m_{\text{LB}} = \tilde{\Omega} \left(\max \left\{ \frac{\sqrt{d}}{\varepsilon^2 \pi_*}, t_{\text{mix}} \right\} \right),$$

where t_{mix}, π_* are as in Theorem 4.1.

Remark 4.4 As the proof shows, for any $0 < \pi_* < 1/(2(d+1))$, a testing problem can be constructed that achieves the $\frac{\sqrt{d}}{\varepsilon^2 \pi_*}$ component of the lower bound. Moreover, for doubly-stochastic $\overline{\mathbf{M}}$, we have $\pi_* = 1/d$, which shows that the upper bound cannot be improved in all parameters simultaneously.

5 Overview of techniques

For both upper and lower bounds, we survey existing techniques, describe their limitations vis-à-vis our problem, and highlight the key technical challenges as well as our solutions for overcoming these.

5.1 Upper bounds

Naïve approach: testing-by-learning. We mention this approach mainly to establish a baseline comparison. Wolfer and Kontorovich (2019) showed that in order to (ε, δ) -learn an unknown d -state ergodic Markov chain $(\mathbf{M}, \boldsymbol{\mu})$ under the $\|\cdot\|$ distance, a single trajectory of length $m_{\text{LEARN}}(\mathbf{M}) = \tilde{O} \left(\frac{1}{\pi_*} \max \left\{ \frac{d}{\varepsilon^2}, t_{\text{mix}} \right\} \right)$ is sufficient. It follows that one can test identity with sample complexity

$$m = \max \{ m_{\text{LEARN}}(\mathbf{M}), m_{\text{LEARN}}(\overline{\mathbf{M}}) \}.$$

This naïve bound, aside from being much looser than bounds provided in the present paper, has the additional drawback of depending on the unknown \mathbf{M} and, in particular, being completely uninformative when the latter is not ergodic.

Reduction to iid testing. Our upper bound in Theorem 4.1 is achieved via the stratagem of invoking an existing iid distribution identity tester as a black box (this is also the general approach of Daskalakis et al. (2018)). Intuitively, given the reference chain $\overline{\mathbf{M}}$, we can compute its stationary distribution $\boldsymbol{\pi}$ and thus know roughly how many visits to expect in each state. Further, computing the mixing time t_{mix} gives us confidence intervals about these expected visits (similar to Wolfer and Kontorovich (2019), via the concentration bounds of Paulin (2015)). Hence, if a chain fails to visit each state a “reasonable” number of times, our tester in Algorithm 1 rejects it. Otherwise, given that state i has been visited as expected, we can apply an iid identity tester to its conditional distribution. The

unknown Markov chain passes the identity test if every state’s conditional distribution passes its corresponding iid test.

A central technical challenge in executing this stratagem is the fact that conditioning on the number of visits creates dependencies on the sample, thereby breaking the Markov property. To get around this difficulty, we use a similar scheme as in Daskalakis et al. (2018). The latter involves introducing an infinite continuation of the sampled trajectory, and subsequently bounding the error of considering this idealized setting, through a covering argument (see Section 6.1.1).

Finally, invoking the tester of Valiant and Valiant (2017) as a black box, it is straightforward to sharpen the worst-case bound in Theorem 4.1 to the instance-specific one in Theorem 4.2.

5.2 Lower bounds

A lower bound of $\Omega(d/\varepsilon^2)$ is immediate via a reduction from the testing problem of Daskalakis et al. (2018) to ours (see Remark 8.1). Although our construction for obtaining the sharper lower bound of $\Omega(\sqrt{d}/(\pi_* \varepsilon^2))$ shares some conceptual features with the constructions in Hao et al. (2018); Wolfer and Kontorovich (2019), a considerably more delicate analysis is required here. Indeed, the technique of tensorizing the KL divergence, instrumental in the lower bound of Wolfer and Kontorovich, would yield (at best) a sub-optimal estimate of $\Omega(1/(\pi_* \varepsilon^2))$ in our setting. Intuitively, bounding TV via KL divergence is too crude for our purposes. Instead, we take the approach of reducing the problem, via a covering argument, to one of iid testing, and construct a family of Markov chains whose structure allows us to recover the Markov property even after conditioning on the number of visits to a certain “special” state. The main contribution for this argument is the decoupling technique of Lemma B.4. The second lower bound is based on the construction of Wolfer and Kontorovich, for which the mixing time and accuracy of the test can both be controlled independently. Curiously, the aforementioned argument cannot be invoked verbatim for our problem, and so we introduce here the twist of considering half-covers of the chains (Lemma B.5), concluding the argument with a two-point technique. This adaptation shaves a logarithmic factor off the corresponding learning problem.

6 Proofs

6.1 Proof of Theorem 4.1

In order to prove Theorem 4.1, we design a testing procedure, describe its algorithm, and further proceed

with its analysis.

Input: $d, \varepsilon, \delta, \overline{\mathbf{M}}, \boldsymbol{\pi}, (X_1, \dots, X_m)$
Output: ACCEPT = 0 or REJECT = 1
 $\mathbf{N} \leftarrow \mathbf{0} \in \mathbb{R}^d$
for $t \leftarrow 1$ **to** $m - 1$ **do**
 | $\mathbf{N}(X_t) \leftarrow \mathbf{N}(X_t) + 1$
end
for $i \leftarrow 1$ **to** d **do**
 | **if** $|\mathbf{N}(i) - (m - 1)\boldsymbol{\pi}(i)| > (m - 1)\boldsymbol{\pi}(i)/2$ **then**
 | | **return** REJECT
 | **end**
end
for $i \leftarrow 1$ **to** d **do**
 | $\mathbf{T} \leftarrow \mathbf{0} \in \mathbb{R}^d$
 | **for** $t \leftarrow 1$ **to** $m - 1$ **do**
 | | $\mathbf{T}(j) \leftarrow \mathbf{T}(j) + \mathbf{1}\{X_t = i\} \mathbf{1}\{X_{t+1} = j\}$
 | **end**
 | **if** $\text{DistIIDIdentityTester}(\varepsilon, d, \overline{\mathbf{M}}(i, \cdot), \mathbf{T})$ is
 | REJECT **then**
 | | **return** REJECT
 | **end**
end
return ACCEPT

Algorithm 1: The testing procedure \mathcal{T}

6.1.1 The testing procedure

For an infinite trajectory X_1, X_2, X_3, \dots drawn from \mathbf{M} , and for any $i \in [d]$, we denote the random hitting times to state i , $\tau_1^{(i)} \doteq \inf \{t \geq 1 : X_t = i\}$, and for $s > 1$, $\tau_s^{(i)} \doteq \inf \{t > \tau_{s-1}^{(i)} : X_t = i\}$. Fixing $\mathbf{n} \in \mathbb{N}^d$ and $i \in [d]$, let us define, following Daskalakis et al. (2018), the mapping

$$\Psi_{\mathbf{n}}^{(i)} : [d]^\infty \rightarrow [d]^{\mathbf{n}(i)}$$

$$(X_1, X_2, X_3, \dots) \mapsto X_{\tau_1^{(i)}+1}, X_{\tau_2^{(i)}+1}, \dots, X_{\tau_{\mathbf{n}(i)}^{(i)}+1},$$

which outputs, for a trajectory drawn from \mathbf{M} , the $\mathbf{n}(i)$ first states that have been observed immediately after hitting i . It is a consequence of the Markov property that the coordinates of $\Psi_{\mathbf{n}}^{(i)}(X_1^\infty)$ be independent and identically distributed according to the conditional distribution defined by the i th state of \mathbf{M} . Namely,

$$\left(X_{\tau_s^{(i)}+1} \right)_{s \in [\mathbf{n}(i)]} \sim \mathbf{M}(i, \cdot)^{\otimes \mathbf{n}(i)}.$$

Remark: For an infinite trajectory, this mapping is *well-defined* almost surely, provided the chain is *irreducible*, while for a finite draw of length m , $\tau_s^{(i)}$ can be infinite for some s , such that proper definition of $\Psi_{\mathbf{n}}^{(i)}$ is a random event that depends on \mathbf{n}, m and the mixing properties of the chain.

For $X_1^m \sim (\mathbf{M}, \boldsymbol{\mu})$, we define our identity tester $\mathcal{T}(\mathbf{X})$ in terms of d sub-testers $\mathcal{T}^{(i)}(\mathbf{X})$, $i \in [d]$, whose definition we defer until further in the analysis. Intuitively, each $\mathcal{T}^{(i)}$ requires at least a “reasonable” number of visits to i , i.e. a lower bound on $\mathbf{n}(i)$.

$$\mathcal{T}(X_1^m) : [d]^m \rightarrow \{0, 1\}$$

$$X_1^m \mapsto 1 - \mathbf{1} \left\{ \begin{array}{l} \forall i \in [d], \Psi_{\mathbf{n}}^{(i)}(X_1^m) \text{ is well-defined} \\ \text{and} \\ \forall i \in [d], \mathcal{T}^{(i)}(\Psi_{\mathbf{n}}^{(i)}(X_1^m)) = 0 \end{array} \right\}$$

6.1.2 Analysis of the tester

Completeness. Consider the two following events

$$\mathcal{E}_\infty \doteq \left\{ \forall i \in [d], \mathcal{T}^{(i)}(\Psi_{\mathbf{n}}^{(i)}(X_1^\infty)) = 0 \right\},$$

$$\mathcal{E}_{\Psi_{\mathbf{n}}} \doteq \left\{ \forall i \in [d], \Psi_{\mathbf{n}}^{(i)}(X_1^m) \text{ is well-defined} \right\}.$$

The probability that the tester correctly outputs 0 for a trajectory sampled from the reference chain $\overline{\mathbf{M}}$ is

$$\mathbf{P}_0(\mathcal{T}(X_1^m) = 0)$$

$$\stackrel{(i)}{=} \mathbf{P}_0 \left(\forall i \in [d], \mathcal{T}^{(i)}(\Psi_{\mathbf{n}}^{(i)}(X_1^m)) = 0 \text{ and } \mathcal{E}_{\Psi_{\mathbf{n}}} \right)$$

$$\stackrel{(ii)}{=} \mathbf{P}_0 \left(\forall i \in [d], \mathcal{T}^{(i)}(\Psi_{\mathbf{n}}^{(i)}(X_1^\infty)) = 0 \text{ and } \mathcal{E}_{\Psi_{\mathbf{n}}} \right)$$

$$\stackrel{(iii)}{=} \mathbf{P}_0(\mathcal{E}_\infty \cap \mathcal{E}_{\Psi_{\mathbf{n}}}) \stackrel{(iv)}{\geq} \mathbf{P}_0(\mathcal{E}_\infty) - \mathbf{P}_0(\neg \mathcal{E}_{\Psi_{\mathbf{n}}}),$$

where (i) is by definition of $\mathcal{T}(X_1^m)$, (ii) stems from the fact that in the event where $\Psi_{\mathbf{n}}^{(i)}(X_1^m)$ is well-defined,

$$\Psi_{\mathbf{n}}^{(i)}(X_1^m) = \Psi_{\mathbf{n}}^{(i)}(X_1^\infty)$$

holds, while (iii) is by definition of \mathcal{E}_∞ , and (iv) is by the following covering argument

$$\mathcal{E}_\infty \subset (\mathcal{E}_\infty \cap \mathcal{E}_{\Psi_{\mathbf{n}}}) \cup \neg \mathcal{E}_{\Psi_{\mathbf{n}}}.$$

Further setting $\mathbf{n} = \frac{1}{2}(m - 1)\boldsymbol{\pi}$, where $\boldsymbol{\pi}$ is the stationary distribution of the reference chain, and from an application of the union bound,

$$\mathbf{P}_0(\neg \mathcal{E}_{\Psi_{\mathbf{n}}}) \leq \sum_{i \in [d]} \mathbf{P}_0 \left(\Psi_{\mathbf{n}}^{(i)}(X_1^m) \text{ is not well-defined} \right)$$

$$\leq \sum_{i \in [d]} \mathbf{P}_0(N_i < \mathbf{n}(i)),$$

$$\leq \sum_{i \in [d]} \mathbf{P}_0(|N_i - \mathbf{E}_\pi[N_i]| > (m - 1)\boldsymbol{\pi}(i)/2),$$

where $N_i \doteq \sum_{t=1}^{m-1} \mathbf{1}\{X_t = i\}$ is the number of visits to state i (not counting the final state at time m). For $m \geq c \frac{t_{\text{mix}}}{\pi_*} \ln \left(\frac{d}{\delta \pi_*} \right)$, $c \in \mathbb{R}_+$,

$$\mathbf{P}_0(\neg \mathcal{E}_{\Psi_{\mathbf{n}}}) \leq \frac{\delta}{3}, \quad (6.1)$$

using the Bernstein-type concentration inequalities of Paulin (2015) as made explicit in Wolfer and Kontorovich (2019, Lemma 5). Observe that no properties of the unknown chain were invoked in this deduction.

We are left with lower bounding $\mathbf{P}_0(\mathcal{E}_\infty)$, the probability that all state-wise testers correctly output 0 in the idealized case where they have access to enough samples. We first recall some standard results (see for example Waggoner, 2015).

Lemma 6.1 (iid $(\varepsilon, 2/5)$ -testing to identity) *Let $\varepsilon > 0$ and $d \in \mathbb{N}$. There exists a universal constant C_{IID} and a tester \mathcal{T}_{FIX} , such that for any reference distribution $\bar{\mathbf{D}} \in \Delta_d$, and any unknown distribution $\mathbf{D} \in \Delta_d$, for a sample of size $m_{\text{IID},2/5} \geq C_{\text{IID}} \frac{\sqrt{d}}{\varepsilon^2}$ drawn iid from \mathbf{D} , \mathcal{T}_{FIX} can distinguish between the cases $\mathbf{D} \equiv \bar{\mathbf{D}}$ and $\|\mathbf{D} - \bar{\mathbf{D}}\|_{\text{TV}} > \varepsilon$ with probability $3/5$.*

Lemma 6.2 (BPP amplification) *Given any $(\varepsilon, 2/5)$ -identity tester \mathcal{T} for the iid case with sample complexity $m_{\text{IID},2/5}$, and any $0 < \delta < 2/5$, we can construct (via a majority vote) an amplified tester \mathcal{T}_{BPP} such that for $m \geq 18 \ln(\frac{2}{\delta}) m_{\text{IID},2/5}$, \mathcal{T}_{BPP} can distinguish the cases $\mathbf{D} \equiv \bar{\mathbf{D}}$ and $\|\mathbf{D} - \bar{\mathbf{D}}\|_{\text{TV}} > \varepsilon$ with confidence $1 - \delta$.*

Since for $i \in [d]$, $\Psi_{\mathbf{n}}^{(i)}(X_1^\infty) \sim \mathbf{M}(i, \cdot)^{\otimes \mathbf{n}(i)}$, and from a union bound, we can invoke Lemma 6.1 and Lemma 6.2. Ensuring that $\forall i \in [d], \mathbf{n}(i) \geq c \frac{\sqrt{d}}{\varepsilon^2} \ln(\frac{1}{\delta})$, $c \in \mathbb{R}_+$, promises that

$$\mathbf{P}_0(\mathcal{E}_\infty) \geq 1 - \delta/3. \quad (6.2)$$

Soundness. In this case, recall that $X_1^m \sim \mathbf{M}$ such that $\exists i_0 \in [d], \|\bar{\mathbf{M}}(i_0, \cdot) - \mathbf{M}(i_0, \cdot)\|_{\text{TV}} > 2\varepsilon$. Following similar arguments as in the completeness case, we upper bound the error probability of the tester

$$\begin{aligned} & \mathbf{P}_1(\mathcal{T}(X_1^m) = 0) \\ &= \mathbf{P}_1\left(\forall i \in [d], \mathcal{T}^{(i)}(\Psi_{\mathbf{n}}^{(i)}(X_1^m)) = 0 \text{ and } \mathcal{E}_{\Psi_{\mathbf{n}}}\right) \\ &= \mathbf{P}_1\left(\forall i \in [d], \mathcal{T}^{(i)}(\Psi_{\mathbf{n}}^{(i)}(X_1^\infty)) = 0 \text{ and } \mathcal{E}_{\Psi_{\mathbf{n}}}\right) \\ &\leq \mathbf{P}_1\left(\forall i \in [d], \mathcal{T}^{(i)}(\Psi_{\mathbf{n}}^{(i)}(X_1^\infty)) = 0\right) \\ &\leq \mathbf{P}_1\left(\mathcal{T}^{(i_0)}(\Psi_{\mathbf{n}}^{(i_0)}(X_1^\infty)) = 0\right). \end{aligned}$$

From Lemma 6.1 and Lemma 6.2, for $\mathbf{n}(i_0) \geq c \frac{\sqrt{d}}{\varepsilon^2} \ln(\frac{1}{\delta})$, it is the case

$$\mathbf{P}_1\left(\mathcal{T}^{(i_0)}(\Psi_{\mathbf{n}}^{(i_0)}(X_1^\infty)) = 0\right) \leq \delta/3. \quad (6.3)$$

Finally, combining (6.1), (6.2) and (6.3) finishes proving the theorem. \square

6.2 Proof of Theorem 4.2

This claim follows immediately from the analysis of the iid instance-optimal tester (Valiant and Valiant, 2017), which is invoked to test the conditional distributions of each state.

Lemma 6.3 (Valiant and Valiant 2017) *Let $\varepsilon > 0$ and $d \in \mathbb{N}$. There exists a universal constant C_{IO} , such that for any reference distribution $\bar{\mathbf{D}} \in \Delta_d$, there exists a tester $\mathcal{T}_{\text{FIX,IO}}$ such that for any unknown distribution $\mathbf{D} \in \Delta_d$, for a sample of size $m_{\text{IO},2/5} \geq C_{\text{IO}} \frac{\|\bar{\mathbf{D}}\|_{2/3}}{\varepsilon^2}$ drawn iid from \mathbf{D} , $\mathcal{T}_{\text{FIX,IO}}$ can distinguish between the cases $\mathbf{D} \equiv \bar{\mathbf{D}}$ and $\|\mathbf{D} - \bar{\mathbf{D}}\|_{\text{TV}} > \varepsilon$ with probability $3/5$.*

We simply have to ensure that for each state i , $\mathbf{n}(i) \geq C_{\text{IO}} \frac{\|\bar{\mathbf{M}}^{i,\cdot}\|_{2/3}}{\varepsilon^2}$, i.e. $\frac{1}{2}(m-1)\pi(i) \geq C_{\text{IO}} \frac{\|\bar{\mathbf{M}}^{i,\cdot}\|_{2/3}}{\varepsilon^2}$, whence the theorem. \square

6.3 Proof of Theorem 4.3

The metric domination result in Lemma 8.1 immediately implies a lower bound of $\Omega(d/\varepsilon^2)$ (see Remark 8.1). We now construct two independent and more delicate lower bounds of $\Omega\left(\frac{\sqrt{d}}{\varepsilon^2 \pi_\star}\right)$ and $\Omega(dt_{\text{mix}})$. Let $\mathcal{M}_{d,t_{\text{mix}},\pi_\star}$ be the collection of all d -state Markov chains whose stationary distribution is minorized by π_\star and whose mixing time is at most t_{mix} . Our goal is to lower bound the minimax risk, defined by

$$\mathcal{R}_m \doteq \inf_{\mathcal{T}} \sup_{\bar{\mathbf{M}}, \mathbf{M}} [\mathbf{P}_0(\mathcal{T} = 1) + \mathbf{P}_1(\mathcal{T} = 0)], \quad (6.4)$$

where the inf is over all testing procedures $\mathcal{T} : (X_1, \dots, X_m) \mapsto \{0, 1\}$, and the sup is over all $\bar{\mathbf{M}}, \mathbf{M} \in \mathcal{M}_{d,t_{\text{mix}},\pi_\star}$ such that $\|\bar{\mathbf{M}} - \mathbf{M}\| > \varepsilon$.

The analysis is simplified by considering $(d+1)$ -state Markov chains with d even; an obvious modification of the proof handles the case of odd d . Fix $0 < p_\star \leq 1/(2(d+1))$ and $0 < \varepsilon < 1/2$, and define $\mathbf{p} \in \Delta_{d+1}$ by

$$\mathbf{p}(d+1) = p_\star \text{ and } \mathbf{p}(i) = (1-p_\star)/d, \quad i \in [d]. \quad (6.5)$$

Define the collection of $(d+1)$ -state Markov chain transitions matrices,

$$\mathcal{G}_{p_\star} \doteq \{\mathbf{M}_\eta : \eta \in \Delta_d\},$$

$$\mathbf{M}_\eta \doteq \begin{pmatrix} \mathbf{p}(1) & \dots & \mathbf{p}(d) & p_\star \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{p}(1) & \dots & \mathbf{p}(d) & p_\star \\ \eta(1) & \dots & \eta(d) & 0 \end{pmatrix}.$$

The stationary distribution $\boldsymbol{\pi}$ of a chain of this family is given by

$$\boldsymbol{\pi}(i) = \frac{\boldsymbol{p}(i) + p_\star \boldsymbol{\eta}(i)}{1 + p_\star}, i \neq d+1, \quad \boldsymbol{\pi}(d+1) = \frac{p_\star}{1 + p_\star},$$

and for $p_\star < \frac{d\boldsymbol{\eta}(i)}{d+1}, \forall i \in [d]$, we have $\pi_\star = \boldsymbol{\pi}(d+1)$. We define the n th hitting time for state i as the random variable $H_i^{(n)} \doteq \inf \left\{ t \geq 1 : \sum_{s=1}^t \mathbf{1}\{X_s = i\} = n \right\}$; in words, this is the first time t at which state i has been visited n times. Suppose that $\mathbf{X} = (X_1, \dots, X_m) \sim (\mathbf{M}, \boldsymbol{p})$ for some $\mathbf{M} \in \mathcal{G}_{p_\star}$. For any $n \in \mathbb{N}$, the n th hitting time $H_{d+1}^{(n)}$ to state $d+1$, stochastically dominates² the random variable $\sum_{s=1}^n R_s$, where each R_s is an independent copy distributed as $\text{Geometric}(p_\star)$. To see this, consider a similar chain where the value 0 in the last row is replaced with p_\star with $\boldsymbol{\eta}$ appropriately re-normalized; clearly, the modification can only make it easier to reach state $d+1$. Continuing, we compute $\mathbf{E}[\sum_{s=1}^n R_s] = n/p_\star$ and $\mathbf{Var}[\sum_{s=1}^n R_s] = n(1-p_\star)/p_\star^2 \leq n/p_\star^2$. The Paley-Zygmund inequality implies that for $m < n/(2p_\star)$,

$$\begin{aligned} \mathbf{P}\left(H_{d+1}^{(n)} > m\right) &\geq \mathbf{P}\left(\sum_{s=1}^n R_s > m\right) \\ &\geq \mathbf{P}\left(\sum_{s=1}^n R_s > \frac{1}{2}\mathbf{E}\left[\sum_{s=1}^n R_s\right]\right) \\ &\geq \left(1 + \frac{4\mathbf{Var}[\sum_{s=1}^n R_s]}{\mathbf{E}[\sum_{s=1}^n R_s]^2}\right)^{-1} \geq 1 - \frac{1}{1+n/4} \geq \frac{1}{5}. \end{aligned} \tag{6.6}$$

Define the random variable $N_{d+1} = \sum_{t=1}^m \mathbf{1}\{X_t = d+1\}$, i.e. the number of visits to state $d+1$, and consider a reference Markov chain $\overline{\mathbf{M}} \doteq \mathbf{M}_{\overline{\boldsymbol{\eta}}} \in \mathcal{G}_{p_\star}$, where $\overline{\boldsymbol{\eta}} = \mathbf{U}_d$ is the $[d]$ -supported uniform distribution. Restricting the problem to a subset of the family \mathcal{G}_{p_\star} satisfying the ε -separation condition only makes it easier for the tester, as does taking any mixture \mathbf{M}_Σ of chains of this class in lieu of the sup in (6.4). More specifically, we choose

$$\mathbf{M}_\Sigma \doteq \frac{1}{2^{d/2}} \sum_{\boldsymbol{\sigma} \in \{-1,1\}^{d/2}} \mathbf{M}_{\boldsymbol{\sigma}},$$

where $\mathbf{M}_{\boldsymbol{\sigma}}(i, \cdot) = \boldsymbol{p}$ for $i \in [d]$, and

$$\begin{aligned} \mathbf{M}_{\boldsymbol{\sigma}}(d+1, \cdot) &= \mathbf{D}_{\boldsymbol{\sigma}} \\ &= \left(\frac{1 + \sigma_1 \varepsilon}{d}, \frac{1 - \sigma_1 \varepsilon}{d}, \dots, \frac{1 + \sigma_{d/2} \varepsilon}{d}, \frac{1 - \sigma_{d/2} \varepsilon}{d}\right). \end{aligned}$$

By construction, for all $\boldsymbol{\sigma} \in \{-1,1\}^{d/2}$, we have $\|\mathbf{M}_{\boldsymbol{\sigma}} - \overline{\mathbf{M}}\| = \|\mathbf{D}_{\boldsymbol{\sigma}} - \mathbf{U}_d\|_1 = \varepsilon$. We start both

² A random variable A stochastically dominates B if $\mathbf{P}(A \geq x) \geq \mathbf{P}(B \geq x)$ for all $x \in \mathbb{R}$.

chains with distribution \boldsymbol{p} , defined in (6.5). Now (6.6) implies that for any $n \in \mathbb{N}$ and $m < n/(2p_\star)$, any testing procedure $\mathcal{T} : (X_1, \dots, X_m) \mapsto \{0,1\}$ verifies

$$\begin{aligned} &\mathbf{P}_{\overline{\mathbf{M}}}(\mathcal{T} = 1) + \mathbf{P}_{\mathbf{M}_\Sigma}(\mathcal{T} = 0) \\ &\geq \frac{1}{5} \left(\mathbf{P}_{\overline{\mathbf{M}}}(\mathcal{T} = 1 | \mathcal{E}_n) + \mathbf{P}_{\mathbf{M}_\Sigma}(\mathcal{T} = 0 | \mathcal{E}_n) \right), \end{aligned}$$

where we wrote $\mathcal{E}_n \doteq \{N_{d+1} \leq n\}$. It follows from Le Cam (2012, Chapter 16, Section 4) that

$$\mathcal{R}_m \geq \frac{1}{5} \left(1 - \|\mathbf{P}_{\mathbf{M}_\Sigma}(\mathbf{X} | \mathcal{E}_n) - \mathbf{P}_{\overline{\mathbf{M}}}(\mathbf{X} | \mathcal{E}_n)\|_{\text{TV}}\right),$$

and so it remains to upper bound a total variation distance. For any $\mathbf{M}, \mathbf{M}' \in \mathcal{G}_{p_\star}$, the statistics of the induced state sequence only differ in the visits to state $d+1$.

At this point, we would like to invoke an iid testing lower bound — but are cautioned against doing so naively, as conditioning on the number of visits to a state breaks the Markov property. Instead, in Lemmas B.2, B.3 and B.4 we develop a decoupling technique, which yields

$$\|\mathbf{P}_{\mathbf{M}_\Sigma}(\mathbf{X} | \mathcal{E}_n) - \mathbf{P}_{\overline{\mathbf{M}}}(\mathbf{X} | \mathcal{E}_n)\|_{\text{TV}} \leq \|\mathbf{D}_\Sigma^{\otimes n} - \mathbf{U}_d^{\otimes n}\|_{\text{TV}}.$$

We shall make use of Paninski (2008, Theorem 4), which states:

$$\|\mathbf{D}_\Sigma^{\otimes n} - \mathbf{U}_d^{\otimes n}\|_{\text{TV}} \leq \sqrt{\exp\left(\frac{n^2 \varepsilon^4}{d}\right)}.$$

It follows that

$$\mathcal{R}_m \geq \frac{1}{5} \left(1 - \frac{1}{2} \sqrt{\exp\left(\frac{n^2 \varepsilon^4}{d}\right) - 1}\right).$$

Finally, for the mixture of chains and parameter regime in question, we have $\boldsymbol{\pi}(d+1) = \pi_\star$ and $p_\star/2 \leq \pi_\star \leq p_\star$, so that for $\delta < 1/10, m < \frac{n}{2p_\star}$ and $n \leq \frac{\sqrt{d}}{\varepsilon^2} \sqrt{\ln\left(\frac{4}{\delta^2}(1-5\delta)^2\right)}$, it follows that $\mathcal{R}_m \geq \delta$.

This implies a lower bound of $m = \Omega\left(\frac{\sqrt{d}}{\varepsilon^2 \pi_\star}\right)$ for the testing problem. \square

The proof of the second lower bound in $\Omega(dt_{\text{mix}})$ is deferred to Section A.

7 Computational complexity

Time complexity of designing a test. In order to determine the size of the sample required in order to reach desired proximity ε and confidence $1 - \delta$, and to run the test algorithm, one must first compute $\boldsymbol{\pi}$

and t_{mix} . These two quantities are relative to the reference Markov chain, for which we assume to have a full description of the transition matrix $\overline{\mathbf{M}}$. Interestingly they can be computed *offline*, only once per test definition, no matter how many observed trajectories we will want to test.

The eigenproblem $\pi \overline{\mathbf{M}} = \pi$ can be solved in time $\tilde{\mathcal{O}}(d^\omega)$, where $2 \leq \omega \leq 2.3728639$ (Le Gall, 2014); this is a method for recovering π .

In the reversible case, it is well known (Levin et al., 2009) that the mixing time of an ergodic Markov chain is controlled by its *absolute spectral gap* $\gamma_\star \doteq 1 - \max\{\lambda_2, |\lambda_d|\}$, where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ is the ordered spectrum of $\overline{\mathbf{M}}$, and minimum stationary probability π_\star :

$$\left(\frac{1}{\gamma_\star} - 1\right) \ln 2 \leq t_{\text{mix}} \leq \frac{\ln(4/\pi_\star)}{\gamma_\star}.$$

The full eigen-decomposition used to obtain π of cost $\tilde{\mathcal{O}}(d^\omega)$ also yields λ_\star .

In the non-reversible case, the relationship between the spectrum and the mixing time is not nearly as straightforward, and it is the *pseudo-spectral gap* (Paulin, 2015), $\gamma_{\text{ps}} \doteq \max_{k \in \mathbb{N}} \{\gamma((\overline{\mathbf{M}}^\dagger)^k \overline{\mathbf{M}}^k)/k\}$, where $\overline{\mathbf{M}}^\dagger$ is the time reversal (Fill, 1991) of $\overline{\mathbf{M}}$, that gives effective control:

$$\frac{1}{2\gamma_{\text{ps}}} \leq t_{\text{mix}} \leq \frac{1}{\gamma_{\text{ps}}} \left(\ln \frac{1}{\pi_\star} + 2 \ln 2 + 1 \right).$$

For any $k \in \mathbb{N}$, we have $\frac{\gamma((\overline{\mathbf{M}}^\dagger)^k \overline{\mathbf{M}}^k)}{k} \leq 1/k$, and so γ_{ps} is computable to within an additive error of η in time $\mathcal{O}(d^3/\eta)$.

The instance-specific test requires computing the d values $\|\overline{\mathbf{M}}(i, \cdot)\|_{2/3}$, which is feasible in time $\mathcal{O}(d^2)$.

Time and space complexity of performing a test. The first step is to verify whether each state has been visited a sufficient amount of times $n_i, i \in [d]$, which can be achieved in time $\mathcal{O}(m+d)$ and space $\mathcal{O}(d)$. The second step is to proceed with the state-wise testing strategy, whose runtime will depend on the black-box iid tester. For instance, the test of Valiant and Valiant (2017) can be performed in $\mathcal{O}(m+d)$ operations. Using this as the iid tester, we get an overall time complexity of $\mathcal{O}(d(d+m))$.

8 Comparison with existing work

Let us compare our distance $\|\cdot\|$ to the one of Daskalakis et al. (2018):

$$\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 1 - \rho\left([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}\right),$$

where $\left([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}\right)(i, j) = \left[\sqrt{\mathbf{M}(i, j)\mathbf{M}'(i, j)}\right]$ and ρ is the spectral radius.

Lemma 8.1 *For all $\mathbf{M}, \mathbf{M}' \in \mathcal{M}_d$, $\|\mathbf{M} - \mathbf{M}'\| \geq 2\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}')$. Conversely, for $d \geq 4$, there exist $\mathbf{M}, \mathbf{M}' \in \mathcal{M}_d$ such that $\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 0$ while $\|\mathbf{M} - \mathbf{M}'\| = 1$.*

Proof: For $\mu, \nu \in \Delta_d$, define the Hellinger distance

$$H^2(\mu, \nu) \doteq \frac{1}{2} \sum_{i=1}^d \left(\sqrt{\mu(i)} - \sqrt{\nu(i)} \right)^2.$$

For arbitrary $\mathbf{M}, \mathbf{M}' \in \mathcal{M}_d$, a standard calculation yields

$$\begin{aligned} \|\mathbf{M} - \mathbf{M}'\| &\geq 2 \max_{i \in [d]} H^2(\mathbf{M}(i, \cdot), \mathbf{M}'(i, \cdot)) \\ &= 2 \max_{i \in [d]} \left(1 - \sum_{j=1}^d \sqrt{\mathbf{M}(i, j)\mathbf{M}'(i, j)} \right) \\ &= 2 \left(1 - \min_{i \in [d]} \sum_{j=1}^d \sqrt{\mathbf{M}(i, j)\mathbf{M}'(i, j)} \right) \\ &\geq 2 \left(1 - \rho([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}) \right) \\ &= 2\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}'), \end{aligned}$$

where the second inequality follows from the Perron-Frobenius theorem (Meyer, 2000, Chapter 8). For the second claim, choose any $\theta \in (-1, 1)$ and put

$$\mathbf{M} = \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 + \theta & 1 - \theta \end{pmatrix},$$

$$\mathbf{M}' = \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}.$$

Since they have an identical connected component, $\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 0$, whereas $\|\mathbf{M} - \mathbf{M}'\| = |\theta|$. \square

Remark 8.1 *It follows that ε -identity testing with respect $\text{Dist}_{\sqrt{\cdot}}(\cdot, \cdot)$ reduces to (2ε) -identity testing with respect to $\|\cdot\|$. In particular, the results of Daskalakis et al. (2018) immediately imply a lower bound of $\Omega(d/\varepsilon^2)$ for our testing problem.*

Acknowledgments

We are thankful to John Lafferty for bringing this problem to our attention and to Nick Gravin for the insightful conversations.

References

- T. W. Anderson and L. A. Goodman. Statistical inference about markov chains. *The Annals of Mathematical Statistics*, pages 89–110, 1957.
- M. Anthony and P. L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge, 1999. ISBN 0-521-57353-X. doi: 10.1017/CBO9780511624216. URL <http://dx.doi.org/10.1017/CBO9780511624216>.
- F. Barsotti, A. Philippe, and P. Rochet. Hypothesis testing for markovian models with random time observations. *Journal of Statistical Planning and Inference*, 173:87–98, 2016.
- M. S. Bartlett. The frequency goodness of fit test for probability chains. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 47, pages 86–95. Cambridge University Press, 1951.
- T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 259–269, 2000. doi: 10.1109/SFCS.2000.892113. URL <https://doi.org/10.1109/SFCS.2000.892113>.
- T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 442–451. IEEE, 2001.
- P. Billingsley. Statistical methods in markov chains. *The Annals of Mathematical Statistics*, pages 12–40, 1961.
- Y. Cherapanamjeri and P. L. Bartlett. Testing symmetric markov chains without hitting. In A. Beygelzimer and D. Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 758–785, Phoenix, USA, 25–28 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v99/cherapanamjeri19a.html>.
- C. Daskalakis, N. Dikkala, and N. Gravin. Testing symmetric markov chains from a single trajectory. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 385–409. PMLR, 06–09 Jul 2018.
- I. Diakonikolas, T. Gouleakis, J. Peebles, and E. Price. Collision-based testers are optimal for uniformity and closeness. *Chicago Journal of Theoretical Computer Science*, to appear, 2019+.
- J. A. Fill. Eigenvalue bounds on convergence to stationarity for nonreversible markov chains, with an application to the exclusion process. *The annals of applied probability*, pages 62–87, 1991.
- O. Goldreich. The uniform distribution is complete with respect to testing identity to a fixed distribution. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 23, page 1, 2016.
- O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 68–75. Springer, 2011.
- Y. Hao, A. Orlitsky, and V. Pichapati. On learning markov chains. In *Advances in Neural Information Processing Systems*, pages 648–657, 2018.
- P. G. Hoel. A test for markoff chains. *Biometrika*, 41(3/4):430–433, 1954.
- D. Kazakos. The Bhattacharyya distance and detection between markov chains. *IEEE Transactions on Information Theory*, 24(6):747–754, 1978.
- A. Kontorovich and I. Pinelis. Exact lower bounds for the agnostic probably-approximately-correct (pac) machine learning model. *Ann. Statist.*, 47(5):2822–2854, 2019. ISSN 0090-5364. doi: 10.1214/18-AOS1766.
- L. Le Cam. *Asymptotic methods in statistical decision theory*. Springer Science & Business Media, 2012.
- F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pages 296–303. ACM, 2014.
- D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times, second edition*. American Mathematical Soc., 2009.
- C. Meyer. *Matrix analysis and applied linear algebra*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000. ISBN 0-89871-454-0. URL <https://doi.org/10.1137/1.9780898719512>. With 1 CD-ROM (Windows, Macintosh and UNIX) and a solutions manual (iv+171 pp.).
- L. Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.
- D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20, 2015.
- G. Valiant and P. Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.

- B. Waggoner. L_p testing and learning of discrete distributions. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 347–356, 2015. doi: 10.1145/2688073.2688095. URL <http://doi.acm.org/10.1145/2688073.2688095>.
- G. Wolfer and A. Kontorovich. Minimax learning of ergodic markov chains. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, volume 98 of *Proceedings of Machine Learning Research*, pages 904–930, Chicago, Illinois, 22–24 Mar 2019. PMLR. URL <http://proceedings.mlr.press/v98/wolfer19a.html>.