

Empirical Analysis of a Stochastic Approximation Approach for Computing Quasi-stationary Distributions

Jose Blanchet, Peter Glynn, and Shuheng Zheng

Abstract. This paper studies a method for estimating the quasi-stationary distribution of various interacting particle processes has been proposed by [6, 5, 8]. This method improved upon existing methods in eigenvector estimation by eliminating the need for explicit transition matrix representation and multiplication. However, this method has no firm theoretical foundation. Our paper analyzes the algorithm by casting it as a stochastic approximation algorithm (Robbins-Monro) [12]. In doing so, we prove its convergence and rate of convergence. Based on this insight, we also give an example where the rate of convergence is very slow. This problem can be alleviated by using an improved version of this algorithm that is given in this paper. Numerical experiments are described that demonstrate the effectiveness of this improved method.

Keywords: Stochastic Approximation, Quasi-stationary Distributions.

1 Introduction

The original motivation for this algorithm came from physicists' need to estimate the quasi-stationary distribution of the contact process [6, 5, 8]. A quasi-stationary distribution can be computed via the left principal eigenvector of the transition matrix (transition rate matrix in the continuous-time setting). The method that has been proposed by these physicists is a heuristic based on manipulation of the Kolmogorov

Jose Blanchet · Shuheng Zheng
Columbia University
e-mail: jose.blanchet@ieor.columbia.edu,
JohnZ622@gmail.com

Peter Glynn
Stanford University
e-mail: glynn@stanford.edu

forward equations. The method works in practice but has no firm proof. This paper recognizes the algorithm as a stochastic approximation algorithm which allows us to prove convergence and sufficient conditions for a Central Limit Theorem. We then give an improved version with variance reduction.

This section reviews the relevant related literature on eigenvector estimations. Sect. 2 reviews some background material to the contact process, quasi-stationary distributions, and the basis for the original heuristic. Sect. 3 goes over the stochastic approximation formulation and sketches the proof of convergence (the full proof will be given in a follow-up journal paper [1]). Sect. 4 gives an improved version of the algorithm. Sect. 5 studies the algorithm adapted for continuous-time Markov chains. Sect. 6 goes over several important numerical experiments.

1.1 Related Literature

Power Method

The power method [9] is very simple. We iterate a sequence x_n by computing

$$x_{n+1} = \frac{x_n^T A}{\|x_n^T A\|}$$

This works for any matrix such that the principal eigenvalue has multiplicity one and strictly largest magnitude. The problem is that for Markov chains with extremely large state space, such as the contact process, it would not be feasible to store and compute in such large dimensions (on the order of 2^n for interacting particle systems).

The variant known as inverse method also suffers from similar problems due to the necessity of matrix multiplication.

1.1.1 Monte Carlo Power Method

The Monte Carlo power method involves a random sampling of the values in the matrix in such a way that a sequence converges to the principal eigenvalue. This method works for any matrix A .

We need to define a Markov chain on the *index* of the matrix A : $1, \dots, n$. Call this Markov chain $\{k_n\}$ where a transition from $k_n = \alpha$ to $k_{n+1} = \beta$ depends on the magnitude of $A_{\alpha\beta}$ in the following way

$$P(k_{n+1} = \beta | k_n = \alpha) = \frac{|A_{\alpha\beta}|}{\sum_{\beta} |A_{\alpha\beta}|}$$

with an arbitrary initial distribution generator \mathbf{h} so that

$$P(k_0 = \alpha) = \frac{|h_\alpha|}{\sum_\alpha |h_\alpha|}$$

Then we define a random sequence of variables W_n recursively:

$$W_0 = \frac{h_{k_0}}{p_{k_0}} \quad W_n = W_{n-1} \frac{A^{k_{n-1}k_n}}{p_{k_{n-1}k_n}}$$

It isn't hard to verify that

$$\lim_{n \rightarrow \infty} \frac{E[W_n \mathbf{f}_{k_n}]}{E[W_{n-1} \mathbf{f}_{k_{n-1}}]} = \lim_{n \rightarrow \infty} \frac{h^T A^n f}{h^T A^{n-1} f} = \lambda_{max}$$

for any \mathbf{f} . This method grows according to $O(Nnm)$, where N is the number of states in your Markov chain, n is the step number at when you terminate $E[W_n f_{b_n}]$, and m is the number of independent Monte Carlo paths that you use to construct $E[W_n f_{b_n}]$. However, in the contact process case, we can reduce this to $O(Knm)$ where K is the number of nodes in the graph. The major drawback to this method is that it will only give you the approximate eigenvalue. In order to get the eigenvector, a lot more work is required especially for large matrices such as ours.

1.1.2 Other Methods

[2] is a good survey of other principal eigenvector estimation algorithms. [13] is a very recent pre-print of a stochastic method that is related but different from our method. The mirror descent method method of [11] is also another alternative.

2 Background and Motivation

2.1 Contact Process

A contact process is a continuous-time Markov chain (CTMC) $(X'_1, \dots, X'_n) \in \{0, 1\}^n$, where $t \geq 0$ is the time, with an associated connected graph (V, E) such that

- (i) $|V| = n$
- (ii) Individual nodes transition from 1 to 0 at an exponential rate of 1
- (iii) Individual nodes transition from 0 to 1 at rate λr where r is the fraction of neighbors that are in state 1

This CTMC has 2^n states. The state $(0, 0, \dots, 0)$ is an absorbing state and the remaining states are all transient.

This CTMC will eventually reach the absorbing state but physicists are interested in the ‘‘pseudo-equilibrium’’ behavior in the long time before absorption happens.

The hindrance of large state space can be alleviated in the rare cases where a compressed representation is possible, such as the case of a contact process on complete graph. In that case, we only need to know the total number of “on” states rather than the identities of all the “on” states.

2.2 Quasi-stationary Distribution

2.2.1 Discrete-Time Version

[3] proposed the concepts of quasi-stationary distribution and quasi-limiting distribution for the discrete-time Markov chains. Assume that 0 is the absorbing state and $1, \dots, n$ are absorbing, we can write the Markov transition matrix as

$$P = \begin{bmatrix} 1 & 0 \\ \alpha & Q \end{bmatrix}$$

First we define the conditional transition probabilities

$$\begin{aligned} d_j^\pi(n) &= P(X_n = j | X_0 \sim \pi, X_1, \dots, X_{n-1} \neq 0) \\ &= \frac{\pi^t Q^{n-1} e_j}{\pi^t Q^{n-1} e} \end{aligned}$$

where $\{e_i\}$ is the standard basis for R^n and e is the vector of all 1's. If there is a distribution π over the transient states such that $d^\pi(n)$ is constant, then we call d^π the quasi-stationary distribution.

Under the assumption that the substochastic matrix Q is irreducible (not necessarily aperiodic), it is straightforward to see that the quasi-stationary distribution exists and is the unique solution to principal eigenvector problem

$$d^t Q = \rho d$$

by the Perron-Frobenius theorem.

Assuming Q is aperiodic and the condition that if $|\rho_2| = |\rho_3|$, we require the multiplicity of ρ_2 to be no less than the multiplicity of ρ_3 , we have that

$$d_j^\pi(n) \rightarrow d_j + O\left(n^k \frac{|\rho_2|}{\rho_1}\right)$$

Note that the rate of convergence depends on the ratio between the second eigenvalue and principle eigenvalue.

2.2.2 Continuous-Time

If we think about the transition rate matrix of a CTMC under similar setup (irreducibility and ergodicity), then it ([4]) can be said that

$$d_j^\pi(t) \rightarrow d_j + o(e^{t(\rho' - \rho_1)})$$

where d is the principal left-eigenvector of the rate matrix corresponding to the transient states with associated eigenvalue ρ_1 . I.e.

$$d^t R = \rho_1 d^t$$

where R is the rate matrix of the CTMC.

2.3 Physicist's Heuristic

Under the setting of a continuous-time Markov chain with rate matrix R and absorbing state 0 (without loss of generality, we can combine all absorbing states into one state), if we define $p_{ij}(t) = P(X_t = j | X_0 = i)$ and $P_{is}(t) = 1 - p_{i0}(t)$, then we have that quasi-stationary distribution $d_j = \lim_{t \rightarrow \infty} \frac{p_{ij}(t)}{P_{is}(t)}$. If we apply the Kolmogorov forward equation (known to physicists as the master equation), we get that

$$\frac{dp_{ij}(t)}{dt} = \sum_k p_{ik} R_{kj} \quad (1)$$

and

$$\frac{dP_{is}(t)}{dt} = \frac{d}{dt}(1 - p_{i0}(t)) = - \sum_k p_{ik} R_{k0} \quad (2)$$

Intuitively by the definition of d_j , we have that $p_{ij}(t) \approx d_j P_{is}(t)$ in the quasi-stationary time window (t large enough). So we can apply this to the preceding two equations and get

$$\begin{aligned} d_j \frac{dP_{is}(t)}{dt} &= \sum_k d_k P_{is}(t) R_{kj} \\ \frac{dP_{is}(t)}{dt} &= - \sum_k d_k P_{is}(t) R_{k0} \end{aligned}$$

Combine the two and get

$$d_j \left(\sum_k d_k R_{k0} \right) + \sum_k d_k R_{kj} = 0$$

This gives us a non-linear equation for the equilibrium condition for the quasi-stationary distribution \mathbf{d} . We can think of this as the stationary point of the forward equation

$$\frac{d(d_j)}{dt} = \sum_k d_k R_{kj} + d_j \left(\sum_k d_k R_{k0} \right) \quad (3)$$

The first part of this equation is the standard Kolmogorov forward equation, while the second part redeposits the probability of hitting the absorbing states onto all the non-absorbing states according to the current distribution d_j .

This suggests the following algorithm

Algorithm 1 Algorithm for estimating quasi-stationary distribution

- (i) Start the Markov chain in a non-absorbing state.
 - (ii) Simulate the Markov chain normally.
 - (iii) If the Markov chain hits the absorbing state, re-sample the starting position based on an empirical estimate of the quasi-stationary distribution up until that point and go to step 2. That is, we sample a non-absorbing state according to a weight proportional to the amount of time that such a state has been visited so far throughout the whole algorithm.
 - (iv) The samples after a large enough time window will be drawn approximately from the quasi-stationary distribution.
-

For large enough time, the dynamics of the Markov chain will be governed by (3), which means we can obtain the quasi-stationary distribution by examining the empirical distribution after some large enough time.

3 Stochastic Approximation Analysis of the Algorithm

In this section, we will re-organize Algorithm 1 into a stochastic approximation algorithm. This will let us rigorously prove a convergence result and Central Limit Theorem for the algorithm.

3.1 Formal Description of the Algorithm

We will now write down a precise description of the above Algorithm 1. Let our state space be the finite set S and $T \subset S$ be the set of transient states. Let μ_n be a probability measure over transient states. μ_n will be the cumulative empirical distribution up until the n -th iteration. Let Q be the substochastic matrix over the transient states, and $\{X_n^l\}_n$ be the l th Markov chain in the simulation, and $\tau^l = \min\{k \geq 0 | X_k^l \notin T\}$ (the hitting time of the absorbing state), we can write our algorithm as

$$\mu_{n+1}(x) = \frac{\left(\sum_{l=0}^n \tau^l \right) \mu_n(x) + \left(\sum_{k=0}^{\tau^{n+1}-1} I(X_k^{n+1} = x | X_0^{n+1} \sim \mu_n) \right)}{\sum_{l=0}^{n+1} \tau^l} \quad \forall x \in T$$

for any arbitrary initial distribution μ_0 .

Robbins-Monro, or stochastic approximation algorithms [12, 10], have the form

$$\mu_{n+1} = \mu_n + \alpha_n Y(\mu_n)$$

where

$$\sum_n \alpha_n = \infty \quad \sum_n \alpha_n^2 < \infty \quad \alpha_n \geq 0 \quad \alpha_n \rightarrow 0$$

and $Y(\cdot)$ is a collection of vector-valued random variables for each possible point in the state-space. Note that over the years, the form of the Robbins-Monro algorithm has been extended. The form here is the classical version.

Under certain conditions, which will be discussed rigorously in [1], μ_n converges to root of the function $g(\mu) \triangleq E[Y(\mu)]$. We will transform μ_n into stochastic approximation

$$\mu_{n+1}(x) = \mu_n(x) + \left(\frac{1}{n+1} \right) \left(\frac{\sum_{l=0}^{\tau^{(n+1)}-1} \left(I(X_l^{(n+1)} = x) - \mu_n(x) \right)}{\frac{1}{n+1} \sum_{j=0}^{n+1} \tau^{(j)}} \right)$$

where

$$Y(\mu) = \frac{\sum_{l=0}^{\tau-1} (I(X_l = x | X_0 \sim \mu) - \mu(x))}{\frac{1}{n+1} \sum_{j=0}^{n+1} \tau^{(j)}}$$

The denominator is problematic because it depends on the whole history of μ_n and not just on the present state. To solve this, we artificially consider another state T_n in the following way.

Stochastic approximation scheme for the main algorithm

$$T_{n+1} = T_n + \frac{1}{n+2} (\tau^{(n+1)} - T_n) \Rightarrow T_n = \frac{1}{n+1} \sum_{j=0}^n \tau^{(j)}$$

$$\mu_{n+1}(x) = \mu_n(x) + \tag{4}$$

$$\left(\frac{1}{n+1} \right) \left(\frac{\sum_{l=0}^{\tau^{(n+1)}-1} \left(I(X_l^{(n+1)} = x | X_0^{(n+1)} \sim \mu_n) - \mu_n(x) \right)}{T_n + \frac{\tau^{(n+1)}}{n+1}} \right) \tag{5}$$

we can therefore define

$$Y_n(\mu_n, T_n)(x) \triangleq \frac{\sum_{l=0}^{\tau^{(n+1)}-1} \left(I(X_l^{(n+1)} = x) - \mu_n(x) \right)}{T_n + \frac{\tau^{(n+1)}}{n+1}}$$

$$Z_n(\mu_n, T_n) \triangleq (\tau^{(n+1)} - T_n)$$

So now we have a stochastic approximation path (μ_n, T_n) , where the control parameters are (μ, T) , that fits into the Robbins-Monro scheme above.

Remark 1 Please note that the iterates μ_n are constrained in $H \triangleq \{\mathbf{x} \in \mathbb{R}_+^n \mid \sum x_i = 1\}$, the $(n-1)$ -dimensional simplex.

We can also define a similar algorithm for the continuous-time Markov chain by keeping track of the amount of time a Markov chain spends in each transient state. This is given in Sect. 5.

We can summarize the conditions for our algorithm in the following theorem taken from Blanchet, Glynn, and Zheng (2012):

Theorem 1 Given an irreducible absorbing Markov chain over a finite state space S of cardinality d , let

- (i) Matrix Q denoting the transition rates over the non-absorbing states
- (ii) Let μ_0 be a probability vector over the non-absorbing states
- (iii) Let $T_0 \in \mathbb{R}^+$

Then there exists a unique quasi-stationary distribution μ satisfying the equations

$$\begin{aligned}\mu^t Q &= \lambda \mu \\ \mu^t \mathbf{1} &= 1 \\ \mu &\geq 0\end{aligned}$$

and the Algorithm 1 converges to the point $(\mu, \frac{1}{1-\lambda})$ with probability 1.

Furthermore, if λ_{PV} is the principal eigenvalue of Q and λ_{NPV} are the other eigenvalues and they satisfy

$$\operatorname{Re} \left(\frac{1}{1 - \lambda_{NPV}} \right) < \frac{1}{2} \left(\frac{1}{1 - \lambda_{PV}} \right) \quad \forall \lambda_{NPV} \text{ non-principal eigenvalues}$$

Then

$$\sqrt{n}(\mu_n - \mu) \rightarrow^d N(0, V)$$

for some matrix V .

3.2 Sketch of Proof of Convergence

Our proof in [1] rests on the use of the ODE method [10] where we are required to examine the asymptotics of the coupled dynamical system

$$\dot{\mu}(t) = E_{\mu(t), T(t)} \left[\frac{\sum_{l=0}^{\tau-1} (I(X_l = \cdot | X_0 \sim \mu)) - \tau \mu(x)}{T(t)} \right]$$

$$\text{(Define)} \quad A \triangleq (I - Q)^{-1} = \frac{1}{T} [\mu^t(t)A - (\mu^t A \mathbf{1})\mu^t(t)]$$

$$\begin{aligned}\dot{T}(t) &= E_\mu[\tau] - T \\ &= \mu^t(t)(I - Q)^{-1}\mathbf{1} - T(t)\end{aligned}$$

where $\mu \in R^n$ and $T \in R^+$.

In [1], we were able to show that for a given initial position in the probability simplex, the solution to the above dynamical system exists and converges to its stationary point which is the unique point that satisfies

$$\begin{aligned}\mu^T Q &= \rho \mu^T \\ \sum \mu_i &= 1 \\ \mu_i &\geq 0\end{aligned}$$

and $\rho = 1 - \frac{1}{E_\mu(\tau)}$.

By Theorem 4.2.1 from [10], we can conclude that μ_n converges to the quasi-stationary distribution for all initial configurations (μ_0, T_0) .

By Theorem 10.2.1 from [10], we conclude that a Central Limit Theorem exists as long as the Jacobian of the ODE vector field has spectral radius less than -0.5 . This is equivalent to requiring that

$$Re\left(\frac{1}{1 - \lambda_{NPV}}\right) < \frac{1}{2} \left(\frac{1}{1 - \lambda_{PV}}\right) \quad \forall \lambda_{NPV} \text{ non-principal eigenvalues} \quad (6)$$

where the λ 's are the eigenvalues of the Q matrix.

4 Variations on the Existing Algorithm with Improved Rate of Convergence

One interesting question to ask is what happens when the sufficient conditions for Central Limit Theorem is not met. We will study a simple example consisting of two states.

4.1 Counter Example to CLT

Imagine we have a Markov chain with three states $\{0, 1, 2\}$ and transition matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ \varepsilon & \frac{1-\varepsilon}{2} & \frac{1-\varepsilon}{2} \\ \varepsilon & \frac{1-\varepsilon}{2} & \frac{1-\varepsilon}{2} \end{bmatrix}$$

Obviously the state $\{0\}$ is the absorbing state. In this setup, because of symmetry, our Algorithm 1 reduces to

- (i) With probability $\frac{1-\varepsilon}{2}$ sample either the state 1 or 2 (without knowing the previous state, this is ok by symmetry)
- (ii) With probability ε , sample from either 1 or 2 according to the empirical distribution up until this point.

We recognize this as a self-interacting Markov chain.

A self-interacting Markov chain [7] is a stochastic process $\{X_n\}$ such that

$$P(X_{n+1} \in dx | \mathcal{F}_n) = \Phi(S_n)(dx)$$

where Φ is a function that transforms one measure into another measure and S_n is the empirical measure generated by $\{X_k\}_{k=0}^n$.

Then our quasi-stationary algorithm reduces to the empirical process of a SIMC X_n governed by

$$P(X_{n+1} = dz | \mathcal{F}_n) = \int K(x, dz) dS_n(dx)$$

where the kernel is given by

$$K(x, dz) = \varepsilon \delta_x(dz) + \left(\frac{1-\varepsilon}{2}\right) [\delta_1(dz) + \delta_2(dz)]$$

The sufficient condition for CLT (6) or this problem translates to requiring $\varepsilon < 0.5$.

When the CLT is violated however, [7] states that over a very general class of bounded and measurable functions f

$$E[(S_n(f) - \bar{S}_n(f))^2] = \Theta\left(\frac{1}{n^{2(1-\varepsilon)}}\right)$$

where $S_n(f) = \int f(x) dS_n(x)$, $\bar{S}_n(f) = E[S_n(f)]$. Although this doesn't technically contradict with the existence of a \sqrt{n} -CLT, it does suggest that the scaling sequence is $n^{1-\varepsilon}$ instead of \sqrt{n} .

4.2 The Parallel Algorithm

There is a variant of the algorithm that can offer significant practical benefits. Imagine that at each iteration, instead of there being one run of the Markov chain until absorption, we have M independent runs. Such that

$$\begin{aligned} \mu_{n+1}(x) &= \frac{\mu_n(x) \left(\sum_{l=0}^n \sum_{m=1}^M \tau^{l,m} \right) + \sum_{m=1}^M \left[\sum_{k=0}^{\tau^{n+1,m}-1} I(X_k^{n+1,m} = x | X_0^{n+1,m} \sim \mu_n) \right]}{\sum_{l=0}^{n+1} \sum_{m=1}^M \tau^{l,m}} \\ &= \mu_n(x) + \frac{1}{n+1} \frac{\sum_{m=1}^M \left[\sum_{k=0}^{\tau^{n+1,m}-1} \left(I(X_k^{n+1,m} = x | X_0^{n+1,m} \sim \mu_n) \right) - \tau^{n+1,m} \mu_n(x) \right]}{\frac{1}{n+1} \sum_{l=0}^{n+1} \sum_{m=1}^M \tau^{l,m}} \end{aligned}$$

Again we have to include an extra dimension

$$\mu_{n+1} = \mu_n + \frac{1}{n+1} \left(\frac{\sum_{m=1}^M \left[\sum_{k=0}^{\tau^{n+1,m}-1} \left(I(X_k^{n+1,m} = \cdot | X_0^{n+1,m} \sim \mu_n) \right) - \tau^{n+1,m} \mu_n \right]}{T_n + \frac{1}{n+1} \sum_{m=1}^M \tau^{n+1,m}} \right)$$

$$T_{n+1} = T_n + \frac{1}{n+2} \left(\sum_{m=1}^M \tau^{m,n+1} - T_n \right) \Rightarrow T_n = \frac{1}{n+1} \sum_{j=0}^n \sum_{m=1}^M \tau^{m,j}$$

After some derivation, we obtain the dynamical system

$$\dot{\mu}(t) = \frac{M}{T(t)} (\mu^t (I - Q)^{-1} - (\mu^t (I - Q)^{-1} \mathbf{1}) \mu^t)$$

$$\dot{T}(t) = M \mu^t (I - Q)^{-1} \mathbf{1} - T$$

Very similarly, we know that

$$\mu(t) \rightarrow \bar{\mu}$$

$$T(t) \rightarrow \frac{M}{1 - \bar{\lambda}}$$

If we let g^μ and g^T denote the dynamical system's components, then we obtain the Jacobian

$$\nabla_\mu g^\mu = \frac{M}{T} \left((I - Q)^{-1} - (I - Q)^{-1} \mathbf{1} \mu^T - \frac{1}{1 - \bar{\lambda}} I \right)$$

$$\nabla_T g^\mu = -\frac{M}{T^2} (\mu^t (I - Q)^{-1} - (\mu^t (I - Q)^{-1} \mathbf{1}) \mu^t)$$

$$\nabla_\mu g^T = M (I - Q)^{-1} \mathbf{1}$$

$$\nabla_T g^T = -1$$

So the condition for which the Central Limit Theorem remains the same:

$$\text{Re} \left(\frac{1}{1 - \lambda_{NPV}} \right) < \left(\frac{1}{2} \right) \frac{1}{1 - \lambda_{PV}} \quad \forall \lambda_{NPV} \text{ non principal eigenvalues}$$

where λ_{PV} is the principal eigenvalue of Q and λ_{NPV} is the non-principal eigenvalue of Q .

Although the Central Limit Theorem does not always hold, the variance of the stochastic approximation noise is lower with bigger M . This means that if we have enough independent Markov chain iterations across different processors, the algorithm would perform better. See Section 5 for empirical performance.

5 Continuous-Time Markov Chains

5.1 Formulation and Convergence

So far, the exposition has assumed that the Markov chain of interest is a discrete-time process. It is straightforward to adapt our method for continuous-time processes (such as the contact process). If we denote the transition rate matrix of the CTMC in the following block form

$$T = \begin{bmatrix} 0 & 0 \\ N & Q \end{bmatrix}$$

then we can write the algorithm as

$$\begin{aligned} \mu_{n+1}(x) &= \frac{\mu_n(x) \left(\sum_{l=0}^n \sum_{m=1}^M \tau^{l,m} \right) + \sum_{m=1}^M \left[\int_0^{\tau^{n+1,m}} I(X^{n+1,m}(s) = x | X_0^{n+1,m} \sim \mu_n) ds \right]}{\sum_{l=0}^{n+1} \sum_{m=1}^M \tau^{l,m}} \\ &= \mu_n(x) + \frac{1}{n+1} \frac{\sum_{m=1}^M \left[\int_0^{\tau^{n+1,m}} \left(I(X^{n+1,m}(s) = x | X_0^{n+1,m} \sim \mu_n) \right) ds - \tau^{n+1,m} \mu_n(x) \right]}{\frac{1}{n+1} \sum_{l=0}^{n+1} \sum_{m=1}^M \tau^{l,m}} \end{aligned}$$

By a similar approach as the discrete-time case, we deduce the related dynamical system

$$\begin{aligned} \dot{\mu}(t) &= -\frac{M}{T(t)} \left(\mu^t Q^{-1} - (\mu^t Q^{-1} \mathbf{1}) \mu^t \right) \\ \dot{T}(t) &= -M \mu^t Q^{-1} \mathbf{1} - T \end{aligned}$$

It is straightforward to adapt the Perron-Frobenius theorem to transition rate matrices such as Q by decomposing $Q = A - bI$ where A is an irreducible matrix. We know the existence of a principal eigenvector of positive entries $\bar{\mu}$ (with eigenvalue smaller than 0) such that

$$\bar{\mu}^t Q = \bar{\lambda} \bar{\mu}^t$$

We can easily check that the stationary point, and with more work the limit point of the dynamical system satisfies

$$\begin{aligned} \bar{\mu}^t Q &= \frac{1}{\bar{\mu}^t Q^{-1} \mathbf{1}} \bar{\mu}^t = \bar{\mu}^t \\ \bar{T} &= -M \bar{\mu}^t Q^{-1} \mathbf{1} = -M \frac{1}{\bar{\lambda}} \end{aligned}$$

Hence we have proven that the CTMC version of the algorithm converges to the quasi-stationary distribution.

5.2 Rate of Convergence

The Jacobian of the dynamical system is given by

$$\begin{aligned}\nabla_{\mu} g^{\mu} &= -\frac{M}{T} (Q^{-1} - Q^{-1} \mathbf{1} \mu^T - (\mu^T Q^{-1} \mathbf{1}) I) \\ \nabla_T g^{\mu} &= \frac{M}{T^2} (\mu^T Q^{-1} - (\mu^T Q^{-1} \mathbf{1}) \mu^T) \\ \nabla_{\mu} g^T &= -M Q^{-1} \mathbf{1} \\ \nabla_T g^T &= -1\end{aligned}$$

When evaluated at the stationary point $(\bar{\mu}, \bar{T})$, we get the matrix

$$\begin{bmatrix} -\bar{\lambda} \left(Q^{-1} - Q^{-1} \mathbf{1} \bar{\mu}^T - \frac{1}{\bar{\lambda}} I \right) & -M Q^{-1} \mathbf{1} \\ \mathbf{0} & -1 \end{bmatrix}$$

If λ_Q is any non-principal eigenvalue of Q , then the sufficient condition for CLT becomes

$$2\lambda_{PV} > \text{Re}(\lambda_Q)$$

5.3 Uniformization

Because these CTMC have finite state space, we can form the associated uniformized Markov chain. Let Q be the transition rate matrix of the non-absorbing states and let $\nu = \max_i (-q_{ii})$, we can form a discrete-time transition matrix

$$\tilde{Q} = I + \frac{1}{\nu} Q$$

It is straightforward to verify that any principal left-eigenvector to Q is also a principal left-eigenvector to \tilde{Q} . Hence we apply the discrete-time algorithm to this DTMC.

6 Numerical Experiments

6.1 Loopy Markov Chain

Let's consider the loopy Markov chain given by the full stochastic matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ \varepsilon & \frac{1-\varepsilon}{2} & \frac{1-\varepsilon}{2} \\ \varepsilon & \frac{1-\varepsilon}{2} & \frac{1-\varepsilon}{2} \end{bmatrix}$$

the eigenvalues of the sub-stochastic matrix are $1 - \varepsilon$ and 0. Hence the sufficient condition for Central Limit Theorem to hold is requiring $\varepsilon < 0.5$. A series of numerical experiments were performed for different values of ε where the L2 error is plotted against time. The observation is summarized in the following table.

Table 1 This table summarizes the behavior of the loopy Markov chain for various ε

ε	CLT Satisfied?	Observation	Figure (in appendix)
0.1	yes	No difference between the performance of different M 's.	A.1
0.4	yes (borderline)	No difference between the performance of different M 's.	A.2
0.6	no (borderline)	Noticeable, but relatively medium difference between small M and larger M . Observed critical $M=2$. Anomalous large error for the $M=10$ run.	A.2
0.98	no	Huge difference between the simulation with small M and larger M . However, some of the simulations with very large M begin to show larger errors than the simulation with medium M 's.	A.4

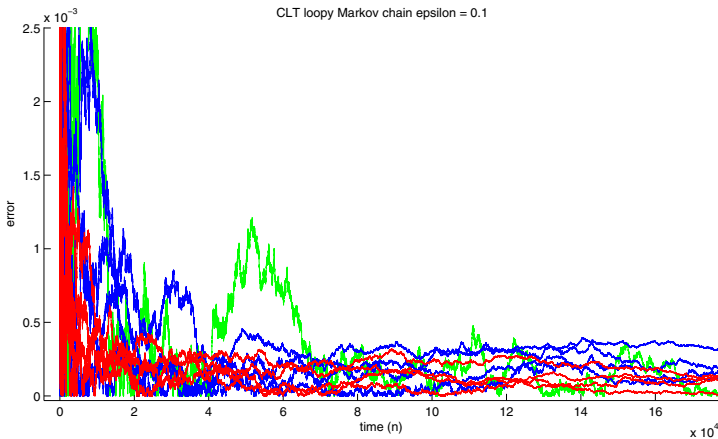


Fig. 1 This figure is the time vs. error plot of the main algorithm ran on a loopy Markov chain with eigenvalues well within the CLT regime ($\varepsilon = 0.1 < 0.5$). Notice the scale of the y-axis. The colors of the different lines represent different runs with different M 's. In this regime which satisfies the CLT for all M , increasing M does not improve the rate of convergence.

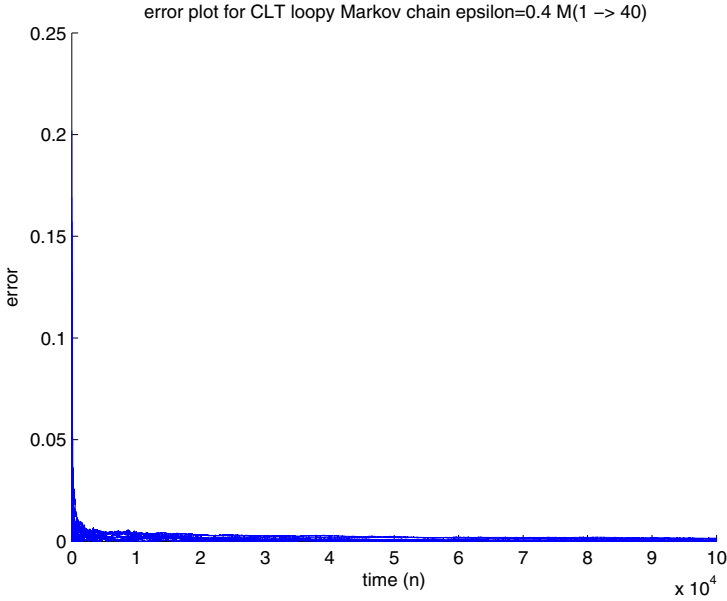


Fig. 2 This figure is the time vs. error plot of the main algorithm run on a loopy Markov chain with eigenvalues just within the CLT regime ($\epsilon = 0.4 < 0.5$). Just like the previous figure, this figure shows that increasing M does not improve the rate of convergence.

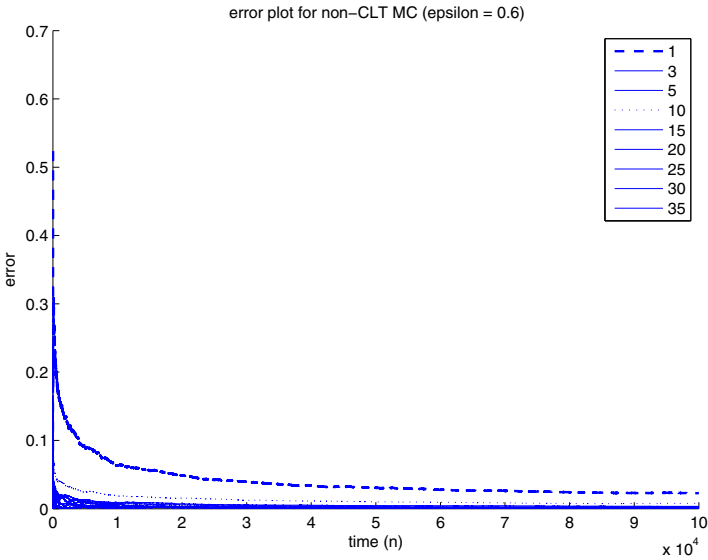


Fig. 3 This is the time vs. error plot of the main algorithm run on a Markov chain with eigenvalues just outside of the CLT regime ($\epsilon = 0.6 > 0.5$). As you can see, there is a noticeable difference between the $M = 1$ simulation and other M simulations. However, there is an anomalous run for $M = 10$. It is probably due to the inherent large variance of the error.

6.2 Contact Process on Complete Graph

6.2.1 Low Infection Rate

Here we apply the algorithm to a special case of the contact process. This is the contact process on a complete graph. This simple case allows the process to be only represented by the number of infected nodes. We picked 10000 nodes and an infection rate of 0.8. The algorithm was run for 1000 iterations. See Fig. 5 for the plot of the estimated distribution vs the true distribution.




Fig. 4 This is the time vs. error plot of the main algorithm run on a Markov chain with eigenvalues just outside of the CLT regime ($\epsilon = 0.98 \gg 0.5$). As you can see, there are huge differences between the $M = 1, 5$ simulation and other M simulations. However, there are anomalous runs for $M = 20, 25$. They are probably due to the inherent large variance of the error.




Fig. 5 This is the time vs. probability plot of the the continuous-time version of the algorithm applied to the contact process on complete graph with 10000 nodes and an infection rate of 0.8.




Fig. 6 This is the time vs. probability plot of he continuous-time version of the algorithm applied to the contact process on complete graph with 10000 nodes and a high infection rate of 1.5.

6.2.2 High Infection Rate

If the infection rate is changed to 1.5, then each iteration of the algorithm would take an extreme long time due to the time it takes to hit the absorbing state. Hence, we uniformized the continuous-time chain to get a discrete-time transition matrix Q . Instead of applying the algorithm to Q , we can apply the algorithm to $0.99 \times Q$ in order to shorten each tour. The algorithm showed high variability on the two different runs. See Fig. 6 for the plot of the estimated distribution vs the true distribution.

7 Discussion and Conclusion

In summary, we have given a rigorous foundation to the algorithm of [5] by recognizing it as a stochastic approximation algorithm. In doing so, we were able to prove its law of large number and fluid limits. A slightly improved algorithm is also proposed and this algorithm significantly improves rate of convergence for some cases.

There also exists a class of projection-based stochastic approximation algorithms $\theta_{n+1} = \Pi[\theta_n + \varepsilon_n Y_n]$ that can be applied to our algorithm. Namely, we can discard the ‘‘T’’ dimension in our algorithm and replace the normalizing denominator by a projection operator. Unfortunately, this algorithm works very poorly in practice.

We have tested our algorithm on countable state space processes such as the $M/M/1/\infty$ queue with success. Proving the convergence of this algorithm in this setting is currently an open problem.

Acknowledgements. Support from the NSF foundation through the grants CMMI-0846816 and CMMI-1069064 is gratefully acknowledged.

References

1. Blanchet, J., Glynn, P., Zheng, S.: Theoretical analysis of a stochastic approximation algorithm for computing quasi-stationary distributions. Technical report, Columbia University (2012)
2. Chatterjee, C.: Adaptive algorithms for first principal eigenvector computation. *Neural Networks* 18, 145–159 (2005)
3. Darroch, J.N., Seneta, E.: On quasi-stationary distributions in absorbing discrete-time finite markov chains. *Journal of Applied Probability* 2(1), 88–100 (1965)
4. Darroch, J.N., Seneta, E.: On quasi-stationary distributions in absorbing continuous-time finite markov chains. *Journal of Applied Probability* 4, 192–196 (1967)
5. de Oliveira, M.M., Dickman, R.: How to simulate the quasistationary state. *Physical Review E* 71 (2005)
6. de Oliveira, M.M., Dickman, R.: Quasi-stationary simulation: the subcritical contact process. *Brazilian Journal of Physics* (2006)
7. del Moral, P., Miclo, L.: Self-interacting markov chains. *Stochastic Analysis and Applications* 24, 615–660 (2006)
8. Dickman, R., Vidigal, R.: Quasi-stationary distributions for stochastic processes with an absorbing state. *Journal of Physics A* 35, 1147–1166 (2002)

9. Golub, G.H., Van Loan, C.F.: Matrix Computations, 3rd edn. Johns Hopkins University Press (1996)
10. Kushner, H.J., Yin, G.: Stochastic approximation and recursive algorithms and applications, vol. 35. Springer (2003)
11. Nazin, A.: Estimating the principal eigenvector of a stochastic matrix: Mirror descent algorithms via game approach with application to pagerank problem. In: 2010 49th IEEE Conference on Decision and Control, pp. 792–797 (2010)
12. Robbins, H., Monro, S.: A stochastic approximation method. *The Annals of Mathematical Statistics* 22, 400–407 (1951)
13. Whiteley, N., Kantas, N.: A particle method for approximating principal eigen-functions and related quantities. arXiv preprint