




Contents lists available at ScienceDirect

Performance Evaluation

journal homepage: www.elsevier.com/locate/peva

Online algorithms for estimating change rates of web pages

Konstantin Avrachenkov^a, Kishor Patil^{a,*}, Gugan Thoppe^b^a INRIA Sophia Antipolis, 2004 Route des Lucioles, Sophia-Antipolis, 06902, France^b Indian Institute of Science, CV Raman Road, Bengaluru, 560012, India

ARTICLE INFO

Article history:

Available online 11 November 2021

Keywords:

Change rate
 Web crawling
 Database synchronization
 Stochastic approximation
 Heavy ball
 Momentum

ABSTRACT

A search engine maintains local copies of different web pages to provide quick search results. This local cache is kept up-to-date by a web crawler that frequently visits these different pages to track changes in them. Ideally, the local copy should be updated as soon as a page changes on the web. However, finite bandwidth availability and server restrictions limit how frequently different pages can be crawled. This brings forth the following optimization problem: maximize the freshness of the local cache subject to the crawling frequencies being within prescribed bounds. While tractable algorithms do exist to solve this problem, these either assume the knowledge of exact page change rates or use inefficient methods such as MLE for estimating the same. We address this issue here.

We provide three novel schemes for online estimation of page change rates, all of which have extremely low running times per iteration. The first is based on the law of large numbers and the second on stochastic approximation. The third is an extension of the second and includes a heavy-ball momentum term. All these schemes only need partial information about the page change process, i.e., they only need to know if the page has changed or not since the last crawled instance. Our main theoretical results concern asymptotic convergence and convergence rates of these three schemes. In fact, our work is the first to show convergence of the original stochastic heavy-ball method when neither the gradient nor the noise variance is uniformly bounded. We also provide some numerical experiments (based on real and synthetic data) to demonstrate the superiority of our proposed estimators over existing ones such as MLE. Our algorithms are readily applicable to the synchronization of databases and network inventory management.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

The world wide web is a highly complex entity: it has a lot of interlinked information, and both the information and the links keep evolving. Nevertheless, even in this challenging setup, one still expects a search engine to provide accurate and up-to-date search results instantaneously. To fulfill this expectation, a search engine maintains a local cache of important web pages, that it updates frequently by using a crawler (also referred to as a web spider or a web robot). Specifically, the job of a crawler [1–5] is (a) to access various pages on the web at specific frequencies so as to determine if any changes have happened to the content since the last crawled instance; and (b) to update the local cache whenever there is a change.¹ There are, however, two key constraints on the different crawling frequencies. The first is due to limitations

* Corresponding author.

E-mail addresses: k.avrachenkov@inria.fr (K. Avrachenkov), kishor88k@gmail.com (K. Patil), gthoppe@iisc.ac.in (G. Thoppe).¹ A web crawler is also supposed to discover new pages, but we do not focus on this task in this work.

on the available bandwidth. The second one, known as the politeness constraint, arises because of the bounds placed by servers on the number of pages that can be accessed in a short amount of time. The search engine thus needs to solve the following optimization problem: maximize the freshness of the local database subject to the crawling frequencies satisfying the above constraints.

In the early 2000s, the web crawling problem used to be formulated as follows [6–8]. The whole web consists of n pages, all have equal importance, and there are no politeness constraints. Further, the times at which different pages change are independent Poisson point processes with different rates [6,9]. On the search engine side, the local cache consists of a copy of each of these n pages. Each copy is updated at regular intervals of time by crawling the original page at a certain (known) frequency. Finally, the freshness of the local cache at time $t \geq 0$ is defined to be $r \in [0, 1]$ if r fraction of the local elements matches the actual versions on the web. The goal then is to find an update policy that maximizes the time-averaged freshness of the local cache and, also, satisfies the bandwidth constraint.

Finding an exact solution to this problem is hard. Hence, numerical solutions were obtained in [6,7] for small values of n . These showed that the optimal crawling policy could be very different to both the uniform as well as the proportional policy, i.e., crawling each page at the same frequency or at one that is proportional to its change rate. In fact, somewhat surprisingly, it was also found that the optimal policy may often include ignoring pages that change too frequently, i.e., not crawling them at all.

In 2003, the freshness definition was modified to include different weights for different pages depending on their importance, e.g., represented as the frequency of requests for different pages [10]. This was done in line with the view that only a finite number of pages can be crawled in any given time frame; hence, to improve the utility of the local database, the freshness criteria should be biased more towards important pages. Numerical solutions, again for small n , confirm that page weights do substantially influence the optimal crawling policy.

While the general n case is still unsolved, a recent breakthrough work [11] showed how an optimal randomized crawling policy can, nonetheless, be found very efficiently (in just $O(n \log n)$ operations). In particular, this solution pertains to the case where, for each web page, even the set of access times forms a Poisson point process. An approach to derandomize this policy to handle the original setup with periodic crawling is also discussed there. In synthetic experiments, this resultant policy is claimed to show performances very similar to the one obtained via numerical solutions. This work was recently extended to cover the case with politeness constraints as well [12].

A separate study [13,14] provides a Whittle index based dynamic programming approach to optimize the schedule of a web crawler. In that approach, the page/catalog freshness estimate also influences the optimal crawling policy.

As can be seen, several algorithms do exist to determine the optimal crawling policy. However, they either presume prior knowledge of the exact page change rates, which is unrealistic in practice, or, alternatively, use inefficient ideas for estimating the same. We now provide a brief overview of such approaches and the issues that plague them.

To the best of our knowledge, three other estimators exist in the literature: the naive estimator [15–17], the Maximum Likelihood Estimator (MLE) [18], and the Moment Matching (MM) estimator [19]. The naive estimator is simply the ratio of the observed number of changes to the total monitoring time period. This is clearly biased since the crawler only has access to partial information about the page change process (remember, it only gets to see if a page has changed or not since the last crawled instance). To overcome this bias issue, MLE instead estimates the rate of change by identifying the parameter value that maximizes the likelihood of the page change observations. This idea performs quite well in experiments; in fact, it also works when access to a page is only possible at irregular intervals of time. However, MLE lacks a closed form expression and suffers from two issues: (a) instability, i.e., the estimator value equals ∞ as long as a page change is detected in every access; and (b) computational intractability, i.e., the estimate needs to be recomputed from scratch each time a new observation is made. The latter makes MLE impractical to use when the data set of observations is quite large. Finally, in the MM estimator, one looks at the fraction of times no changes were detected during page accesses and then, using a moment matching method, estimates the change rate. Unfortunately, like MLE, the MM estimator also suffers from instability and computational issues.

Our exact problem statement and the main contributions can now be summarized as follows. We consider a single page and, as in [11], presume that the page change times and page access times are independent homogeneous Poisson point processes. Thus, each of these processes can be characterized by a single parameter, which we denote here by Δ and p , respectively. Importantly, we assume that only p is known. We then develop three approaches for online estimation of Δ , which only need to know if this page has changed or not between two successive accesses. The key word here is ‘online’. This means, unlike MLE and the MM estimator, our estimates can be incrementally updated using extremely simple, low cost formulas as and when a new observation becomes available. Thus, our estimators do not face computational issues of the kind mentioned above. Also, they do not face any instability issues.

Our first estimator uses the Law of Large Numbers (LLN), while the second and third estimators are based on Stochastic Approximation (SA) principles. Specifically, the update rule for the first estimator is derived using a formula for the probability that there is a page change between two successive accesses. In contrast, the second estimator is constructed via a standard trick in SA. A key ingredient there is a function that is carefully chosen so that it satisfies two properties: (a) noisy estimates of its value for any given input can be easily obtained; and (b) its expected value is linear and, importantly, Δ is its unique zero. The update rule for the third estimator is similar to that of the second one, except that it has an additional momentum term (in the heavy-ball sense). As we show in Section 3.5, it is also possible to view our second and third estimators as a Stochastic Gradient Descent (SGD) method and as an SGD method with heavy-ball momentum,

respectively. Even though we present our results in the context of web crawling, our algorithms are equally applicable to the synchronization of databases [7] and the problem of network inventory management [20].

Our main theoretical result is that all our estimators almost surely (a.s.) converge to Δ ; thus, they all are asymptotically consistent. As far as we know, our result concerning the third estimator is the first to show convergence of an SGD method with heavy-ball momentum when neither the gradient nor the noise variance is a priori assumed to be uniformly bounded. While similar settings have also been dealt with in [21], the analysis there concerns the stochastic analogue of a modified heavy-ball method and not the original one that was proposed in [22]. Separately, we also derive the convergence rates of the first two estimators in the expected error sense. Based on the existing literature, we also provide a loose guess on the convergence rate of the third estimator.

We also provide numerical simulations to compare the performance of our online schemes to each other and also to that of the (offline) MLE estimator. From these experiments, it can be explicitly seen that our estimators give performances comparable to that of MLE. This was a bit surprise to us since our estimators, compared to MLE, have extremely low running times per iteration. Also, unlike MLE, they ignore the actual lengths of intervals between two page accesses. Among our three estimators, LLN and SAM show similar performances and both typically outperform our SA estimator. In particular, the momentum in the third estimator helps in accelerating the estimation whenever $p \ll \Delta$ (the rate at which the page is accessed is much smaller than the rate at which it changes). Our experiments are based on both real (Wikipedia traces) as well as synthetic data sets. In the experiment using Wikipedia traces, we also verify our modeling assumption that the page change process is a Poisson point process.

The rest of this paper is organized as follows. The next section provides a formal summary of this work in terms of the setup, goals, and key contributions. It also gives explicit update rules for all of our online schemes. In Section 3, we formally analyze their convergence and the rates of convergence. The numerical experiments discussed above are given in Section 4. Then, in Section 5, we provide some motivation on how one can use our estimates to find the optimal crawling rates. Finally, we conclude in Section 6 with some future directions.

We note that a shorter version (Avrachenkov et al. 2020, [23]) of this paper appeared in the proceedings of ValueTools 2020 conference. The novel contributions here include (i.) an additional change rate estimation scheme (this is a stochastic approximation scheme with momentum) and its analysis and (ii.) additional experiments including one that compares the performance of all our estimators based on real data (Wikitraces).

2. Setup, goal, and key contributions

The three topics are individually described below.

Setup: Without loss of generality, we work with a single web page. We presume that the actual times at which this page changes is a time-homogeneous Poisson point process in $[0, \infty)$ with a constant but unknown rate Δ . Independently of everything else, this page is crawled (accessed) at the random instances $\{t_k\}_{k \geq 0} \subset [0, \infty)$, where $t_0 = 0$ and the inter-arrival times, i.e., $\{t_k - t_{k-1}\}_{k \geq 1}$, are IID exponential random variables with a known rate p . Thus, the times at which this page is crawled is also a time-homogeneous Poisson point process but with rate p . At time instance t_k , we get to know if the page got modified or not in the interval $(t_{k-1}, t_k]$, i.e., we can access the value of the indicator

$$I_k := \begin{cases} 1, & \text{if the page got modified in } (t_{k-1}, t_k], \\ 0, & \text{otherwise.} \end{cases}$$

The above assumptions are standard in the crawling literature. Nevertheless, we now provide a short justification for the same. Our assumption that the page change process is a Poisson point process is based on the experimental evidence collected in [6,9,24]. An additional validation is provided by us in this work. Specifically, we selected an arbitrary page from the list of frequently edited Wikipedia pages. We extracted the complete history of this web page (exact dates and times of different changes) for a period of five months (April 01, 2020 to August 31, 2020). Thereafter, we calculated the time between successive changes and then used this data to produce a Q-Q plot. This plot confirms that the set of quantiles for the actual data indeed matches linearly with the quantiles of exponential distribution, as predicted. Further details about this experiment can be found in Section 4. Some generalized models for the page change process have also been considered in the literature [8,25]; however, we do not pursue them here.

Our assumption on $\{I_k\}$ is based on the fact that a crawler can only access incomplete knowledge about the page change process. In particular, a crawler does not know when and how many times a page has changed between two crawling instances. Instead, all it can track is the status of a page at each crawling instance and know if it has changed or not with respect to the previous access. Sometimes, it is possible to also know the time at which the page was last modified [2,18], but we do not consider this case here.

Goal: Develop online algorithms for estimating Δ in the above setup. The motivation for doing this is that such estimates can then be used to estimate the optimal crawling rates [11,26]; see Section 5 for more details on this.

Key Contributions: We provide three online methods for estimating the page change rate Δ . The first is based on the law of large numbers, while the second and third are based on stochastic approximation theory, with the third one having an additional momentum component. If $\{x_k\}$, $\{y_k\}$, and $\{z_k\}$, denote the iterates of these three methods, respectively, then their update rules are as shown below.

- *LLN Estimator*: Its k th estimate is given by

$$x_k = p\hat{I}_k / (k + \alpha_k - \hat{I}_k), \quad k \geq 1. \tag{1}$$

Here, $\hat{I}_k = \sum_{j=1}^k I_j$; hence, $\hat{I}_k = \hat{I}_{k-1} + I_k$. Further, $\{\alpha_k\}$ is any positive sequence satisfying the conditions in [Theorem 1](#); e.g., α_k could be $\log k$, \sqrt{k} , or identically 1.

- *SA Estimator*: Given some initial value y_0 , the update rule for the SA estimator is

$$y_{k+1} = y_k + \eta_k [I_{k+1}(y_k + p) - y_k], \quad k \geq 0. \tag{2}$$

Here, $\{\eta_k\}$ is any stepsize sequence that satisfies the conditions in [Theorem 2](#). For example, η_k could be $1/(k + 1)^\eta$ for some constant $\eta \in (0, 1]$.

- *SAM Estimator (SA Estimator with Momentum)*: Given some initial values z_0, z_{-1} , the SAM estimator satisfies

$$z_{k+1} = z_k + \eta_k [I_{k+1}(z_k + p) - z_k] + \zeta_k (z_k - z_{k-1}), \quad k \geq 0. \tag{3}$$

Here, $\{\eta_k\}$ and $\{\zeta_k\}$ are any stepsize sequences that satisfy the conditions given in [Theorem 3](#). For example, one could pick a $\beta \in (1/2, 1]$ and let $\beta_k = 1/(k + 1)^\beta$. Then, $\{\eta_k\}$ and $\{\zeta_k\}$ could be $\{1/(k + 1)^\eta\}$ and $\{(\beta_k - \omega\eta_k)/(\beta_{k-1})\}$, respectively, where $\omega > 0$ is some constant and $\beta + 1/2 < \eta \leq 2\beta$. While we do not show it, we conjecture that one can also pick $\beta \in (0, 1/2]$ and then choose η so that $\beta < \eta \leq 2\beta$. Finally, note that if $\beta = \eta$ and $\omega = 1$, then the asymptotic behavior of (3) will resemble that of (2); this is because $\zeta_k \equiv 0$ then.

We call these methods online because the estimates can be updated on the fly as and when a new observation I_k becomes available. This contrasts the MLE estimator in which one needs to start the calculation from scratch each time a new data point arrives. Also, unlike MLE, our estimators are never unstable; see [Section 3.4](#) for the details.

Our main results include the following. We show that all our three estimators, i.e., x_k, y_k , and z_k , converge to Δ a.s. Further, we show that

1. $\mathbb{E}|x_k - \Delta| = O(\max\{k^{-1/2}, \alpha_k/k\})$, and
2. $\mathbb{E}|y_k - \Delta| = O(k^{-\eta/2})$ if $\eta_k = (k + 1)^\eta$ with $\eta \in (0, 1)$.

Separately, based on existing literature [[27–29](#)], we conjecture that $\mathbb{E}|z_k - \Delta| = \tilde{O}(k^{-\beta/2})$, where \tilde{O} hides logarithmic terms. We also provide several numerical experiments based on real as well as synthetic data for judging the strength of our three proposed estimators.

3. Analysis of the proposed online estimators

Here, we formally discuss the convergence and convergence rates of our three estimators. Thereafter, we compare their behaviors with those that already exist in the literature—the Naive estimator, MLE, and the MM estimator. We end with a summary of existing results on stochastic momentum methods and a discussion on how our convergence result for the SAM estimator extends our current understanding of such methods.

3.1. LLN estimator

Our first aim here is to obtain a formula for $\mathbb{E}[I_1]$. We shall use this later to motivate the form of our LLN estimator.

Let $\tau_1 = t_1 - t_0 = t_1$, where the second equality holds since $t_0 = 0$. Then, as per our assumptions in [Section 2](#), τ_1 is an exponential random variable with rate p . Also, $\mathbb{E}[I_1 | \tau_1 = \tau] = 1 - \exp(-\Delta\tau)$. Hence,

$$\mathbb{E}[I_1] = \Delta / (\Delta + p). \tag{4}$$

This gives the desired formula for $\mathbb{E}[I_1]$.

From this latter calculation, we have

$$\Delta = p\mathbb{E}[I_1] / (1 - \mathbb{E}[I_1]). \tag{5}$$

Separately, because $\{I_k\}$ is an IID sequence and $\mathbb{E}[I_1] \leq 1$, it follows from the strong law of large numbers that $\mathbb{E}[I_1] = \lim_{k \rightarrow \infty} \sum_{j=1}^k I_j / k$ a.s. Thus,

$$\Delta = p \frac{\lim_{k \rightarrow \infty} \sum_{j=1}^k I_j / k}{1 - \lim_{k \rightarrow \infty} \sum_{j=1}^k I_j / k} \quad \text{a.s.}$$

Consequently, a natural estimator for Δ is

$$x'_k = p \frac{\sum_{j=1}^k I_j / k}{1 - \sum_{j=1}^k I_j / k} = p \frac{\hat{I}_k}{k - \hat{I}_k}, \tag{6}$$

where \hat{I}_k is as defined below (1).

Unfortunately, the above estimator faces an instability issue, i.e., $x'_k = \infty$ when I_1, \dots, I_k are all 1. To fix this, one can add a non-zero term in the denominator. The different choices then give rise to the LLN estimator defined in (1).

The following result discusses the convergence and convergence rate of this estimator.

Theorem 1. Consider the estimator given in (1) for some positive sequence $\{\alpha_k\}$.

1. If $\lim_{k \rightarrow \infty} \alpha_k/k = 0$, then $\lim_{k \rightarrow \infty} x_k = \Delta$ a.s.
2. Additionally, if $\lim_{k \rightarrow \infty} \log(k/\alpha_k)/k = 0$, then

$$\mathbb{E}|x_k - \Delta| = O(\max\{k^{-1/2}, \alpha_k/k\}).$$

Proof. Let $\mu = \mathbb{E}[I_1]$, $\bar{I}_k = \hat{I}_k/k$, and $\bar{\alpha}_k = \alpha_k/k$. Then, observe that (1) can be rewritten as $x_k = p\bar{I}_k/(1 + \bar{\alpha}_k - \bar{I}_k)$. Now, $\lim_{k \rightarrow \infty} \bar{I}_k = \mu$ a.s. and $\lim_{k \rightarrow \infty} \bar{\alpha}_k = 0$; the first claim holds due to the strong law of large numbers, while the second one is true due to our assumption. Statement 1. is now easy to see.

We now derive Statement 2. From (5), we have

$$|x_k - \Delta| = \left| x_k - p \frac{\mu}{1 - \mu} \right| \leq p(A_k + B_k),$$

where

$$A_k = \left| \frac{\bar{I}_k}{\bar{\alpha}_k + 1 - \bar{I}_k} - \frac{\mu}{\bar{\alpha}_k + 1 - \mu} \right| \quad \text{and} \quad B_k = \left| \frac{\mu}{\bar{\alpha}_k + 1 - \mu} - \frac{\mu}{1 - \mu} \right|.$$

Since $\alpha_k > 0$ and, hence, $\bar{\alpha}_k > 0$, it follows that

$$B_k = \bar{\alpha}_k \frac{\mu}{(1 - \mu)(\bar{\alpha}_k + (1 - \mu))} \leq \bar{\alpha}_k \frac{\mu}{(1 - \mu)^2}.$$

Similarly,

$$A_k \leq \left(\frac{1 + \bar{\alpha}_k}{1 - \mu} \right) \left(\frac{|\bar{I}_k - \mu|}{\bar{\alpha}_k + 1 - \bar{I}_k} \right).$$

It is now easy to see that $\mathbb{E}[B_k] = O(\bar{\alpha}_k)$. The rest of our arguments concern how fast $\mathbb{E}[A_k]$ decays to 0.

Let $\{\delta_k\}$ be a deterministic sequence that is both non-negative and decays to 0. We will describe how to pick this later. Let k be such that $(1 + \delta_k)\mu < 1$. Then,

$$\mathbb{E} \left[\frac{|\bar{I}_k - \mu|}{\bar{\alpha}_k + 1 - \bar{I}_k} \right] \leq \mathbb{E}[C_k] + \mathbb{E}[D_k],$$

where

$$C_k = \frac{|\bar{I}_k - \mu|}{\bar{\alpha}_k + 1 - \bar{I}_k} 1_{\{\bar{I}_k - \mu \leq \delta_k \mu\}},$$

and

$$D_k = \frac{|\bar{I}_k - \mu|}{\bar{\alpha}_k + 1 - \bar{I}_k} 1_{\{\bar{I}_k - \mu \geq \delta_k \mu\}}.$$

On the one hand,

$$\mathbb{E}[C_k] \leq \frac{\mathbb{E}|\bar{I}_k - \mu|}{\bar{\alpha}_k + 1 - (1 + \delta_k)\mu} \leq \frac{\sqrt{\text{Var}[I_1]}}{\sqrt{k}(\bar{\alpha}_k + 1 - (1 + \delta_k)\mu)}.$$

On the other hand, since $|\bar{I}_k - \mu| \leq 2$ and $1 - \bar{I}_k \geq 0$, it follows by applying the Chernoff bound that

$$\mathbb{E}[D_k] \leq \frac{2}{\bar{\alpha}_k} \Pr\{\bar{I}_k \geq (1 + \delta_k)\mu\} \leq \frac{2}{\bar{\alpha}_k} \exp(-k\delta_k^2\mu/3).$$

Now, pick $\{\delta_k\}$ so that $\delta_k^2 = 6 \log(1/\bar{\alpha}_k)/(k\mu) \vee 0$ for all $k \geq 1$. Notice that this choice is both non-negative and decays to 0 due to our assumptions on $\{\alpha_k\}$; thus, this is a valid choice. It is now easy to see that $\mathbb{E}[C_k] = O(1/\sqrt{k})$ and $\mathbb{E}[D_k] = O(\bar{\alpha}_k)$.

The desired result now follows. \square

3.2. SA estimator

Let I denote a random variable with the same distribution as I_1 . Also, for $y \in \mathbb{R}$, let $H(y, I) = I(y + p) - y$. Next, define $h : \mathbb{R} \rightarrow \mathbb{R}$ using $h(y) := \mathbb{E}[H(y, I)]$. Clearly, $h(y) = p(\Delta - y)/(\Delta + p)$; further, Δ is its unique zero. The theory of stochastic approximation then suggests using the update rule given in (2) for estimating Δ . For later use, also define

$$\begin{aligned} M_{k+1} &= [I_{k+1}(y_k + p) - y_k] - h(y_k) \\ &= \left[I_{k+1} - \frac{\Delta}{\Delta + p} \right] (y_k + p). \end{aligned} \tag{7}$$

We now discuss the convergence and convergence rate of (2).

Theorem 2. Consider the estimator given in (2) for some positive stepsize sequence $\{\eta_k\}$.

1. Suppose that $\sum_{k=0}^{\infty} \eta_k = \infty$ and $\sum_{k=0}^{\infty} \eta_k^2 < \infty$. Then, $\lim_{k \rightarrow \infty} y_k = \Delta$ a.s.
2. Suppose that $\eta_k = 1/(k + 1)^\eta$ for some constant $\eta \in (0, 1)$. Then,

$$\mathbb{E}|y_k - \Delta| = O(k^{-\eta/2}).$$

Proof. For $k \geq 0$, consider the σ -field $\mathcal{F}_k := \sigma(y_j, I_j, j \leq k)$. Then, from (4) and the fact that $\{I_k\}$ is an IID sequence, we get

$$\mathbb{E}[I_{k+1}(y_k + p) - y_k | \mathcal{F}_k] = \frac{\Delta}{\Delta + p} (y_k + p) - y_k = h(y_k).$$

Hence, one can rewrite (2) as

$$y_{k+1} = y_k + \eta_k [h(y_k) + M_{k+1}], \tag{8}$$

where M_{k+1} is as in (7).

Since $\mathbb{E}[M_{k+1} | \mathcal{F}_k] = 0$ for all $k \geq 0$, $\{M_k\}$ is a martingale difference sequence. Consequently, (8) is a classical SA algorithm whose limiting ODE is

$$\dot{y}(t) = h(y(t)). \tag{9}$$

We now make use of Theorem 9 given in the Appendix to establish Statement 1. Accordingly, we verify the four conditions listed there. The stepsize Condition (i.) directly holds due to our assumptions on $\{\eta_k\}$. With regard to Condition (ii.), recall we have already established above that $\{M_k\}$ is a martingale difference sequence with respect to $\{\mathcal{F}_k\}$. The square-integrability condition holds since $|M_{k+1}| \leq |y_k| + p$ which, in turn, implies that $\mathbb{E}[|M_{k+1}|^2 | \mathcal{F}_k] \leq 2(p^2 \vee 1)(1 + |y_k|^2)$, as desired. Next, due to linearity, h is trivially Lipschitz continuous. Further, $h(y) = 0$ if and only if $y = \Delta$. This shows that Δ is the unique equilibrium point of (9). Now, because the coefficient of y in $h(y)$ is negative, it also follows that Δ is the unique globally asymptotically stable equilibrium of (9). This verifies Condition (iii.). We finally consider Condition (iv.) Let $h_\infty(y) := -yp/(\Delta + p)$. Then, clearly, $h_c \rightarrow h_\infty$ uniformly on compacts as $c \rightarrow \infty$. Furthermore, since the coefficient of y is negative in the definition of h_∞ , it is easy to see that the origin is the unique globally asymptotically stable equilibrium of the ODE $\dot{y}(t) = h_\infty(y(t))$, as required. Statement 1. now follows.

We now sketch a proof for Statement 2. First, note that

$$y_{k+1} - \Delta = (1 - a\eta_k)(y_k - \Delta) + \eta_k M_{k+1},$$

where $a = p/(\Delta + p)$. Now, since $\mathbb{E}[M_{k+1} | \mathcal{F}_k] = 0$, we have

$$\mathbb{E}[(y_{k+1} - \Delta)^2 | \mathcal{F}_k] = (1 - a\eta_k)^2 (y_k - \Delta)^2 + \eta_k^2 \mathbb{E}[M_{k+1}^2 | \mathcal{F}_k].$$

Recall that $\mathbb{E}[M_{k+1}^2 | \mathcal{F}_k] \leq C(1 + y_k^2)$ for some constant $C \geq 0$. By substituting this above and then repeating all the steps from the proof of [30, Theorem 3.1], it is not difficult to see that Statement 2 holds as well. \square

3.3. SA estimator with momentum

As stated before, our SAM estimator is the SA estimator discussed above with an additional heavy-ball momentum term. Simulations in Section 4 show that this simple modification results in a drastic improvement in performance.

We now discuss the convergence of the SAM estimator under the assumption that, for $k \geq 0$,

$$\zeta_k = \frac{\beta_k - \omega\eta_k}{\beta_{k-1}}, \tag{10}$$

where $\omega > 0$ is some constant and $\{\beta_k\}$ is some positive real sequence. By substituting (10) and letting $u_k = (z_k - z_{k-1})/\beta_{k-1}$, observe that the update rule in (3) can be rewritten as

$$u_{k+1} = u_k + \gamma_k [I_{k+1}(z_k + p_i) - z_k] - \omega\gamma_k u_k,$$

where $\gamma_k := \eta_k/\beta_k$.

For $k \geq 0$, let M_{k+1} be as in (7). Also, let \mathcal{F}_k denote the σ -field $\sigma(z_0, u_0, I_1, \dots, I_k)$. Clearly, $u_k, z_k \in \mathcal{F}_k$ and $\mathbb{E}[M_{k+1}|\mathcal{F}_k] = 0$. Hence, $\{M_k\}$ is again a martingale difference sequence with respect to the filtration $\{\mathcal{F}_k\}$. Furthermore, since $|M_{k+1}| \leq |z_k| + p$, we have

$$\mathbb{E}[|M_{k+1}|^2|\mathcal{F}_k] \leq 2(p^2 \vee 1)(1 + |z_k|^2). \tag{11}$$

As before, let $a = p/(\Delta + p)$. Also, let $b = \Delta p/(\Delta + p)$ and $\epsilon_k = u_{k+1} - u_k$ for $k \geq 0$. It is then easy to see that one can write down (3) in terms of the following two update rules:

$$u_{k+1} = u_k + \gamma_k[h(u_k, z_k) + M_{k+1}] \tag{12}$$

$$z_{k+1} = z_k + \beta_k[g(u_k, z_k) + \epsilon_k], \tag{13}$$

where $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ are the linear functions given by

$$h(u, z) = b - \omega u - az \quad \text{and} \quad g(u, z) = u.$$

Theorem 3. Consider the SAM estimator given in (3) with ζ_k of the form given in (10). Then $z_k \rightarrow \Delta$ a.s., if one of the following conditions holds true.

1. One-timescale : $\sum_{k \geq 0} \beta_k = \infty$, $\sum_{k \geq 0} \beta_k^2 < \infty$, and $\beta_k = \gamma_k$.
2. Two-timescale: $\sum_{k \geq 0} \beta_k = \sum_{k \geq 0} \gamma_k = \infty$, $\sum_{k \geq 0} (\beta_k^2 + \gamma_k^2) < \infty$, and $\lim_{k \rightarrow \infty} \frac{\beta_k}{\gamma_k} = 0$.

In both these cases, recall that $\gamma_k = \eta_k/\beta_k$.

We state a few remarks concerning this result before discussing its proof.

Remark 4. Examples of $\{\eta_k\}$ and $\{\beta_k\}$ sequences such that the above conditions are satisfied include the following.

- One-timescale: $\beta_k = 1/(k + 1)^\beta$ with $\beta \in (1/2, 1]$ and $\eta_k = 1/(k + 1)^\eta$ with $\eta = 2\beta$.
- Two-timescale: $\beta_k = 1/(k + 1)^\beta$ with $\beta \in (1/2, 1]$ and $\eta_k = 1/(k + 1)^\eta$ with $\frac{1}{2} + \beta < \eta < 2\beta$.

In either case, note that $\lim_{k \rightarrow \infty} \zeta_k = 1$.

Remark 5. The justification for the names given above for the two sets of conditions is as follows. Under the first set of conditions, the update rules in (12) and (13) indeed behave like a one-timescale stochastic approximation algorithm, i.e., both u_k and z_k move on the same timescale. On the other hand, under the second set of conditions, (12) and (13), it behaves like a two-timescale stochastic approximation algorithm. This is because β_k decays to 0 at a much faster rate than γ_k , in turn implying that the changes in $\{z_k\}$, i.e., $\{z_{k+1} - z_k\}$ are of a smaller magnitude than that in $\{u_k\}$.

Remark 6. In the spirit of the above remark, a natural question to consider is the following. Can one pick $\{\eta_k\}$ and $\{\beta_k\}$ so that $\eta_k/\beta_k^2 \rightarrow 0$ or, equivalently, $\gamma_k/\beta_k \rightarrow 0$? That is, can one pick the stepsizes so that u_k now becomes the slowly moving update relative to z_k ? The answer to this question seems to be no. This is because a couple of sufficient conditions needed to guarantee convergence (e.g., Condition (iii.) and (iv.) in Theorem 11) would no longer hold true in this new setup. Furthermore, simulations seem to suggest that the iterates, in fact, race to infinity.

Remark 7. Another question to consider is the following. Can one pick ω , $\{\beta_k\}$, and $\{\eta_k\}$ so that $\zeta_k \rightarrow \zeta$, where ζ is a constant in $(0, 1)$? In particular, can one choose $\omega = (1 - \zeta)$, $\beta_k = 1/(k + 1)^\beta$ with $\beta \in (1/2, 1]$ and then pick $\eta_k = 1/(k + 1)^\beta$ (i.e., $\eta = \beta$) so that $\zeta_k \rightarrow \zeta$? The answer to this second question does not seem to be clear. This is because $\lim_{k \rightarrow \infty} \gamma_k$ would then equal 1. Consequently, again, one of the sufficient conditions to guarantee convergence (e.g., condition (i.) of Theorem 11) would no longer hold. However, simulations in this case do show some promise.

Remark 8. Based on the existing literature on convergence rates for one-timescale and two-timescale linear stochastic approximation [27–30], one can conjecture that $\mathbb{E}|z_k - \Delta| = \tilde{O}(k^{-\beta/2})$ when $\{\beta_k\}$ and $\{\eta_k\}$ are chosen as described in Remark 4. This implies the optimal convergence rate would then again be $\tilde{O}(1/\sqrt{k})$, which matches the bound we have obtained in Theorem 2 for the SA estimator. However, it is possible that this bound may not be tight in the case of the SAM estimator. The is because (13) lacks the martingale difference term and, typically, these are the kind of terms that dictate the convergence rates. Furthermore, simulations in Section 4 suggest that the SAM estimator always converges much faster than the SA estimator.

Proof of Theorem 3. We discuss the two cases one by one.

One-timescale Setup: In this case, the update rules given in (12) and (13) together form a one-timescale stochastic approximation algorithm. More specifically, if we let $v_k = \begin{bmatrix} u_k \\ z_k \end{bmatrix}$, then it follows that

$$v_{k+1} = v_k + \beta_k \left(H(v_k) + \begin{bmatrix} 0 \\ \epsilon_k \end{bmatrix} + \begin{bmatrix} M_{k+1} \\ 0 \end{bmatrix} \right), \tag{14}$$

where $H : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the function defined by

$$H(v) = \begin{bmatrix} b \\ 0 \end{bmatrix} - \begin{bmatrix} \omega & a \\ -1 & 0 \end{bmatrix} v.$$

We now verify the four conditions listed in [Theorem 9](#) and then make use of [Proposition 10](#) (both given in the [Appendix](#)) to show that $v_k \rightarrow \begin{bmatrix} 0 \\ \Delta \end{bmatrix} =: v_*$ a.s. This automatically implies $z_k \rightarrow \Delta$ a.s., which is what we need to prove.

Notice that the stepsize in [\(14\)](#) is β_k . Condition (i.), therefore, trivially holds due to the assumptions made in [Statement 1](#). Next, observe that the martingale difference term in [\(14\)](#) is the vector $\begin{bmatrix} M_{k+1} \\ 0 \end{bmatrix}$. This, along with [\(11\)](#) and the statements above it, shows that Condition (ii.) is true as well.

With regard to Condition (iii.), first note that H is trivially Lipschitz continuous due to the linearity of both its component functions. Next, since $\Delta = b/a$, we have that $H(v) = 0$ if and only if $v = v_*$. Furthermore, since a and ω are strictly positive, the real parts of the eigenvalues of the matrix in the definition of H are also positive. This can be seen from the following set of observations. To begin with, the associated characteristic equation of this matrix is

$$\lambda^2 - \lambda\omega + a = 0.$$

Hence, the roots are $\lambda = (\omega \pm \sqrt{\omega^2 - 4a})/2$. If $\omega^2 < 4a$, then the roots are complex valued; therefore, the real part of both these roots is $\omega/2$ which is clearly positive. On the other hand, if $\omega^2 \geq 4a$, then both the roots are real; further, the smallest of the two roots, i.e., $(\omega - \sqrt{\omega^2 - 4a})/2$, is strictly positive since $a > 0$. This shows that the negative of the matrix given in the definition of H is Hurwitz. Together, these observations show that v_* is the unique globally asymptotically stable equilibrium of the ODE $\dot{v}(t) = H(v(t))$. This verifies Condition (iii.).

Finally, let

$$H_\infty(v) = - \begin{bmatrix} \omega & a \\ -1 & 0 \end{bmatrix} v.$$

Then, it is easy to see that $H_c(v) \rightarrow H_\infty(v)$ uniformly on compact sets as $c \rightarrow \infty$. Also, $H_\infty(v) = 0$ if and only if $v = 0$. Furthermore, as shown before, the negative of the matrix in the definition of H_∞ is Hurwitz. This implies that the origin is the unique globally asymptotically stable equilibrium of the ODE $\dot{v}(t) = H_\infty(v)$. This verifies condition (iv.).

It now remains to check if $\{\epsilon_k\}$ has the decaying behavior described in [Proposition 10](#). Towards this, since $|M_{k+1}| \leq (p + |z_k|)$, we have

$$\left\| \begin{bmatrix} 0 \\ \epsilon_k \end{bmatrix} \right\| \leq C' \gamma_k (1 + |u_k| + |z_k|) \leq C \gamma_k (1 + \|v_k\|)$$

for some constants $C, C' \geq 0$. Now, because γ_k decays to 0 as $k \rightarrow \infty$ due to the assumption in [Statement 1](#)., it follows that $\{\epsilon_k\}$ indeed has the desired behavior.

This completes the proof in the one-timescale setup.

Two-timescale Setup: Since $\beta_k/\gamma_k \rightarrow 0$, one can perceive u_k to be changing on a faster timescale relative to y_k . Hence, the update rules in [\(12\)](#) and [\(13\)](#) can be viewed as a two-timescale stochastic approximation. We now verify the conditions listed in [Theorem 11](#) and then use [Proposition 12](#) (both given in the [Appendix](#)) to conclude $z_k \rightarrow \Delta$ a.s.

Conditions (i.) and (ii.) trivially hold. Hence, we only focus on verifying Conditions (iii.) and (iv.). Because of linearity, h and g are trivially Lipschitz continuous. Next, let $\phi(z) = (b - az)/\omega$ for $z \in \mathbb{R}$. Clearly, ϕ is linear in z and, hence, Lipschitz continuous. Also, $h(\phi(z), z) = 0$. This, along with the fact that the sign in front of u in $h(u, z)$ is negative, shows that $\phi(z)$ is indeed the unique globally asymptotically stable equilibrium of the ODE $\dot{u}(t) = h(u(t), z)$. Next, observe that the ODE $\dot{z}(t) = g(\phi(z(t)), z(t))$ has the form $\dot{z}(t) = (b - az(t))/\omega$. Clearly, this ODE has Δ as its unique globally asymptotically stable equilibrium. This completes the verification of Condition (iii.).

With regard to Condition (iv.), first let h_∞ be the function defined by $h_\infty(u, z) = -\omega u - az$. Also, for $z \in \mathbb{R}$, let $\phi_\infty(z) = -az/\omega$. This function is linear in z and, hence, Lipschitz; also, $\phi_\infty(0) = 0$. Then, on the one hand, $h_c \rightarrow h_\infty$ uniformly on compacts as $c \rightarrow \infty$ and, on the other hand, the ODE $\dot{u}(t) = h_\infty(u(t), z) = -\omega u(t) - az$ indeed has $\phi_\infty(z)$ as its unique globally asymptotically stable equilibrium. Finally, for $z \in \mathbb{R}$, let $g_\infty(z) = -az/\omega$. Then, trivially, $g_c \rightarrow g_\infty$ uniformly on compacts, as $c \rightarrow \infty$. Further, $\dot{z}(t) = g_\infty(z(t)) = -az(t)/\omega$ which indeed has the origin as its unique globally asymptotically stable equilibrium. With this, we finish with verifying Condition (iv.).

Now, as per [Proposition 12](#), we need to show that $\{\epsilon_k\}$ is asymptotically negligible. However, this is indeed true since $|M_{k+1}| \leq (z_k + p)$ which implies $|\epsilon_k| \leq C \gamma_k (1 + |u_k| + |z_k|)$ for some constant $C \geq 0$, and since $\gamma_k \rightarrow 0$.

This shows that $(u_k, z_k) \rightarrow (\phi(\Delta), \Delta) = (0, \Delta)$ a.s., as desired. \square

3.4. Comparison with existing estimators

As far as we know, there are three other approaches in the literature for estimating page change rates—the Naive estimator, MLE, and the MM estimator. The details about the first two estimators can be found in [\[18\]](#) while, for the third

one, one can look at [19]. We now do a comparison, within the context of our setup, between these estimators and the ones that we have proposed.

The Naive estimator simply uses the average number of changes detected to approximate the rate at which a page changes. That is, if $\{q_k\}$ denotes the iterates of the Naive estimator then, in our setup, $q_k = p\hat{I}_k/k$, where \hat{I}_k is as defined below (1). The intuition behind this is the following. If τ_1 is as defined at the beginning of Section 3.1, then

$$\mathbb{E}[N(\tau_1)] = \Delta/p. \quad (15)$$

Thus, the Naive estimator tries to approximate $\mathbb{E}[N(\tau_1)]$ with \hat{I}_k/k then use (15) to determine the change rate.

Clearly, $\mathbb{E}[q_k] = p\Delta/(\Delta + p) \neq \Delta$. Also, from the strong law of large numbers, $q_k \xrightarrow{a.s.} p\Delta/(\Delta + p) \neq \Delta$. Thus, this estimator is not consistent and is also biased. This is to be expected since this estimator does not account for all the changes that occur between two consecutive accesses.

Next, we look at the MLE estimator. Informally, this estimator identifies the parameter value that has the highest probability of producing the observed set of observations. In our setup, the value of the MLE estimator is obtained by solving the following equation for Δ :

$$\sum_{j=1}^k I_j \tau_j / (\exp(\Delta \tau_j) - 1) = \sum_{j=1}^k (1 - I_j) \tau_j, \quad (16)$$

where $\tau_k = t_k - t_{k-1}$ and $\{t_k\}$ is as defined in Section 2. The derivation of this relation is given in [18, Appendix C]. As mentioned in [18, Section 4], the above estimator is consistent.

Note that the MLE estimator makes actual use of the inter-arrival crawl times $\{\tau_k\}$ unlike our two estimators and also the Naive estimator. In this sense, it fully accounts for the information available from the crawling process. Due to this, as we shall see in the experiments section, the quality of the estimate obtained via MLE improves rapidly in comparison to the Naive estimator as the sample size increases.

However, MLE suffers in two aspects: computational tractability and mathematical instability. Specifically, note that the MLE estimator lacks a closed form expression. Therefore, one has to solve (16) by using numerical methods such as the Newton–Raphson method, Fisher’s Scoring Method, etc. Unfortunately, using these ideas to solve (16) takes more and more time as the number of samples grow. Also note that, under the above solution ideas, the MLE estimator works in an offline fashion. In that, each time we get a new observation, (16) needs to be solved afresh. This is because there is no easy way to efficiently reuse the calculations from one iteration into the next (note that the defining Eq. (16) changes in a significant and nontrivial way from one iteration to the other).

Besides the complexity, the MLE estimator is also unstable in two situations. One, when no changes have been detected ($I_j = 0, \forall k \in \{1, \dots, k\}$), and the other, when all the accesses detect a change ($I_j = 1, \forall k \in \{1, \dots, k\}$). In the first setting, no solution exists; in the second setting, the solution is ∞ . One simple strategy to avoid these instability issues is to clip the estimate to some pre-defined range whenever one of bad observation instances occur.

Finally, let us discuss the MM estimator. Here, one looks at the fraction of times no changes were detected during page accesses and then, using a moment matching method, tries to approximate the actual page change rate. In our context, the value of this estimator is obtained by solving $\sum_{j=1}^k (1 - I_j) = \sum_{j=1}^k e^{-\Delta \tau_j}$ for Δ . The details of this equation are given in [19, Section 4]. While the MM idea is indeed simpler than MLE, the associated estimation process continues to suffer from similar instability and computational issues like the ones discussed above.

We emphasize that none of our estimators suffer from any of the issues mentioned above. In particular, all of our estimators are online and have a significantly simple update rule; thus, improving the estimate whenever a new data point arrives is extremely easy. Moreover, all of them are stable, i.e., the estimated values will almost surely be finite. More importantly, the performance of our estimators is comparable to that of MLE. This can be seen from the numerical experiments in Section 4.

3.5. Comparison of Theorem 3 with the literature on stochastic momentum methods

We first provide an alternative characterization of (3). Let $f(z) = \frac{1}{2a}(az - b)^2$, where a and b are as defined below (11), and let h be as defined in Section 3.2. Then, clearly, $h(z) = -\nabla f(z)$. Thus, (3) can be rewritten as

$$z_{k+1} = z_k + \eta_k[-\nabla f(z_k) + M_{k+1}] + \zeta_k(z_k - z_{k-1}),$$

where M_{k+1} is as defined in (7). Consequently, it follows that (3) can also be viewed as an SGD method with a heavy-ball momentum term (similarly, (2) is also an SGD method, but we will not focus on that here).

The above viewpoint now brings forth an interesting question “How does Theorem 3 compare with the existing results on stochastic heavy-ball method and the stochastic variant of Nesterov’s accelerated gradient method?”

While there are numerous results on stochastic momentum methods, surprisingly, most of them hold only under extremely restrictive assumptions: they either need

1. that the gradient of the objective function be uniformly bounded [31,32], or
2. that the noise sequence, i.e., $\{M_{n+1}\}$, be independent of the iterates [33] or, alternatively, its variance be uniformly bounded [31–38].

In our setup, in contrast, the objective function f is quadratic; hence, the magnitude of its gradient grows to infinity as $|z| \rightarrow \infty$. Also, $\mathbb{E}[|M_{k+1}|^2 | \mathcal{F}_k] = (z_k + p)^2 \mathbb{E}[I_{k+1} - \Delta/(\Delta + p)]^2$, which implies that $\mathbb{E}|M_{k+1}|^2 = \frac{\Delta p}{(\Delta + p)^2} \mathbb{E}(z_k + p)^2$. One can thus see that the above assumptions do not directly hold in our case.

To the best of our knowledge, [21,39] are the only other works that similarly do not need the above assumptions. The results in [39], however, only apply to the setup with constant stepsizes. In that case, it is shown there that the iterates converge to a neighborhood of the desired solution but not to the solution itself. On the other hand, [21] does discuss results on convergence and convergence rates of the stochastic heavy-ball method. The analysis there, though, does not apply to the stochastic variant of the original heavy-ball method, i.e., the one proposed in [22, (9)]; instead, it applies to a different variant.

The paper [40] is one other work on stochastic momentum methods that has recently generated significant attention. However, the results there concern a setup where the objective function is of a different nature to the one we consider here. In particular, instead of the gradient, it is assumed there that the objective function itself is defined via an expectation.

In this sense, our work is the first to analyze the stochastic heavy-ball method (in its original form) without a priori presuming that the above two conditions hold. As a matter of fact, it is proved in [21] that the variant which is considered there cannot be analyzed using the standard ODE based stochastic approximation techniques such as the one proposed in [41, Chapter 6]. Our analysis, in contrast, is able to directly make use of the standard approach.

4. Numerical results

We now demonstrate the strength of our estimators using three different experiments. The first one involves real data based on Wikipedia traces. It serves two of our goals. First, we use this experiment to validate our model assumption that the page change process is a stationary Poisson point process. Second, we use it to demonstrate that the estimation quality of our online estimators is comparable to that of the offline MLE estimator. In the second experiment, using synthetic data, we study the impact of Δ and p on our three estimators. In the third experiment, we similarly study how the choices of $\{\alpha_k\}$, $\{\eta_k\}$ and $\{\beta_k\}$ influence the performance. Finally, based on the outcomes of these experiments, we provide some guidelines on which estimator to use in practice.

4.1. Performance on real data (expt. 1)


As mentioned before, our goal here is provide a validation for our model as well as to compare the performance of the different estimators on real data.

To generate the data set, we used Wikipedia traces which are openly available on the web. In particular, we selected an arbitrary page from the list of frequently edited pages on Wikipedia. The title of the page we chose was “Template talk: Did you know”. Next, we extracted the timestamps at which this page was edited over a period of five months (April 01, 2020 to August 31, 2020). We found that this page had changed 4043 times during this period. From the available history, we then calculated the inter-update times of the page change process. The average of these values turned out to be $\Delta = 1.1098$.


Using a Q-Q plot, we then compared the distribution (specifically quantiles) of the collected data to that of an exponential distribution with this Δ rate. The result is given in Fig. 1(a). Notice that the points roughly fall on a straight line. Importantly, this line is very close to the 45° diagonal. This implies that both the sets of quantiles come from the same distribution, thereby confirming that the collected inter-update times indeed follow an exponential distribution whose rate is close to Δ . Equivalently, this implies that the update times come from a Poisson point process with rate close to Δ .

Having verified our assumption, we now compare five different page rate estimators: Naive, MLE, LLN, SA, and SAM. Their performances are given in Figs. 1(b) and 1(c).


The procedure we adopted to obtain these plots was as follows. (Unless specified, we follow the notations from Section 2). Recall that we had access to the actual timestamps at which this Wikipedia page was changed. Keeping this in mind, we artificially generated the crawl instances of this page. These times were sampled from a Poisson point process with rate $p = 0.5$ for Fig. 1(b) and with $p = 0.1$ for Fig. 1(c). We then checked if the page had changed or not between each of the successive crawling instances. This then generated the values of the indicator sequence $\{I_k\}$. For $p = 0.5$, the length of this sequence was 1723 while, for $p = 0.1$, this length turned out to be 340. Using these I_k , p , and inter-update time lengths, we then used the five different estimators mentioned above to find Δ . This gave rise to the trajectories shown in Figs. 1(b) and 1(c). Note that the depicted trajectories correspond to exactly one run of each estimator. The trajectory of the estimates obtained by the SA estimator is labeled Δ^{SA} , etc. The stepsizes chosen for our different estimators are as follows. For our LLN estimator, we had set $\alpha_k \equiv 1$ and, for the SA estimator, we had used $\eta_k = (k + 1)^{-\eta}$ with $\eta = 0.75$. In case of the SAM estimator, we had set $\beta_k = (k + 1)^{-\beta}$ with $\beta = 0.6$ and $\eta_k = (k + 1)^{-\eta}$ with $\eta = 1.2$. (Recall



(a) Q-Q Plot: Real Data versus exponential distribution



(b) $\Delta = 1.10, p = 0.5$



(c) $\Delta = 1.10, p = 0.1$

Fig. 1. Different Estimators: Real Data.

that, in the SAM estimator, the main stepsize is η_k while the stepsize multiplying the momentum term has the form $\zeta_k = (\beta_k - \omega\eta_k)/\beta_{k-1}$.

We now summarize our findings. In Fig. 1(b), we observe that performances of the MLE, LLN, SA and SAM estimators are comparable to each other and all of them outperform the Naive estimator. This last observation is not at all surprising since the Naive estimator completely ignores the changes missed between two successive crawling instances. In contrast

to this, we observe that the estimators behave somewhat differently in Fig. 1(c). Recall that the crawling frequency here is 0.1, which is quite small compared with the value 0.5 that was chosen before. We notice that SAM and MLE estimators perform better than SA and LLN estimators in this scenario.

4.2. Comparison of estimation quality using synthetic data (expt. 2)

Throughout this experiment, we work with synthetic data.

4.2.1. Sample variance and root mean squared error

Our goal here is to study the sample variance and root mean squared error of the estimates obtained from multiple runs of the different estimators. The output is given in Fig. 2.

The data for this experiment is generated as follows. We sample points from two different stationary Poisson point processes, one with parameter $\Delta = 5$ and the other with parameter $p = 3$. We treat the samples from the first process as the times at which an imaginary page changes, and the samples from the second process as the times at which this page is crawled. We then check if the page has changed or not between two successive page accesses. This information is then used to generate the values of the indicator sequence $\{I_k\}$.

We now give $\{I_k\}$, p , as well as the inter-access lengths as input to the five different estimators mentioned before. The stepsizes we use are as follows. For our LLN estimator, we set $\alpha_k \equiv 1$; for the SA estimator, we use $\eta_k = (k + 1)^{-\eta}$ with $\eta = 0.75$; and, for the SAM estimator, we choose $\zeta_k = (\beta_k - \omega\eta_k)/\beta_{k-1}$, where $\eta_k = (k + 1)^{-\eta}$ with $\eta = 1.3$, $\omega = 1$, and $\beta_k = (k + 1)^{-\beta}$ with $\beta = 0.75$. Fig. 2(a) depicts one single run of each of the five estimators.

In Figs. 2(b) and 2(c), the parameter values are exactly the same as in Fig. 2(a). However, we now run the simulation 100 times; the page change times and the page access times are generated afresh in each run. Fig. 2(b) depicts the 95% confidence interval of the obtained estimates, whereas Fig. 2(c) shows the root mean squared value of the difference between the estimated value and actual change rate of the page.

We now summarize our findings. Clearly, in each case, we observe that performances of the MLE, LLN, SA and SAM estimators are comparable to each other and all of them outperform the Naive estimator. The fact that the estimates from our approaches are close to that of the MLE estimator was indeed quite surprising to us. This is because, unlike MLE, our estimators completely ignore the actual lengths of the intervals between two accesses. Instead, they use p , which only accounts for the mean interval length. Note that the root mean square error of the first few samples for MLE is very high (hence, it is not depicted in Fig. 2(c)). This is due to the instability that MLE faces; see Section 3.4. Fig. 2(c) shows that the error in the MLE estimate decays faster as compared to others. We believe this is because the MLE also uses the actual interval lengths in its computation; thus, it uses more information about the crawling process than the other estimators.

While the plots do not show this, we once again draw attention to the fact that the time taken by each iteration in MLE rapidly grows as k increases. In contrast, our estimators take roughly the same amount of time for each iteration.

4.2.2. Impact of Δ and p on performance


In the previous experiments, recall that our different estimators more or less behaved similarly. Our goal now is to vary the values of Δ and p and see if there are any major differences that crop up in their performances. Alongside, we also wish to see the usefulness of the momentum term used in the SAM estimator. The performances in two such interesting scenarios are shown in Figs. 3 and 4. Note that we no longer consider MLE on account of their impractical run times when the $\{I_k\}$ sequence lengths are large.

In Fig. 3, $\Delta = 500$ and $p = 3$, which means the crawling frequency is quite low compared to the frequency at which the page is updated. On the other hand, in Fig. 4, $\Delta = 500$ and $p = 50$; thus, the crawling frequency now is relatively higher. The stepsizes for our different estimators are as follows. For the LLN estimator, we chose $\alpha_k \equiv 1$; for the SA estimator, we chose $\eta_k = (k + 1)^{-\eta}$ with $\eta = 0.8$; and, for the SAM estimator, we chose η_k as before, $\omega = 1$ and $\beta_k = (k + 1)^{-\beta}$ with $\beta = 0.5$ (note that our stepsize choice for the SAM estimator violates the conditions of Theorem 3, but it satisfies the one we made in the conjecture below (3)).


Figs. 3(a) and 4(a) show one single trajectory of our estimators in the two scenarios. We observe that the LLN and SAM estimators perform quite well as compared to the SA estimator in both the scenarios; however, the latter catches up when the p value becomes higher. The impact of the momentum term can also be clearly seen in the low frequency crawling case. In this scenario, note that the crawler will more or less always detects a change. That is, the $\{I_k\}$ sequence will mostly consists of all 1s. In turn, this means that the SA estimator's update rule will almost always have the form $y_{k+1} = y_k + \eta_k p$.

We then ran the simulation 100 times and obtained a plot of the 95% confidence interval and the root mean squared error of our different estimators in the two scenarios. This is shown in Fig. 3(b), 3(c), 4(b), and 4(c). We observe that variance for SA is relatively very low. This is because the SA estimator does not deviate too much from the update rule mentioned in the previous paragraph. The disadvantage, however, is that its estimates typically are quite far away from the actual change rate. Furthermore, this error decreases quite slowly. Another interesting observation from Figs. 3(b) and 3(c) is that the variance of LLN estimator is larger than that of SAM estimator, however, its error decays at much faster rate than that of the SAM estimator.


Compared to Fig. 3, notice that in Fig. 4 that performance of all our estimators improve. However, as shown in Fig. 4(b), the SAM estimator is quite volatile now. Separately, the zoomed-in plot in 4(c) shows that the average error for the SAM estimator drops quite rapidly compared to others in the initial few iterations. However, this advantage disappears after 400 iterations; then on the LLN estimator performs much better.



(a) Performance of single trajectories




(b) Mean estimate (solid) and 95% Confidence interval




(c) Root mean square error


Fig. 2. Synthetic data: $\Delta = 5$, $p = 3$.



(a) Performance of single trajectories




(b) Mean estimate (solid) and 95% Confidence interval




(c) Root mean square error


Fig. 3. Synthetic data: $\Delta = 500$, $p = 3$.



(a) Performance of single trajectories




(b) Mean estimate (solid) and 95% Confidence interval




(c) Root mean square error


Fig. 4. Synthetic data: $\Delta = 500$, $p = 50$.




(a) LLN estimator for different $\{\alpha_k\}$ choices



(b) SA estimator with $\eta_k = (k + 1)^{-\eta}$ for different η choices



(c) SAM estimator with $\eta_k = k^{-0.8}$ for different $\{\beta_k\}$ choices



(d) SAM estimator with $\beta_k = k^{-0.6}$ for different $\{\eta_k\}$ choices

Fig. 5. Impact of $\{\alpha_k\}$, $\{\eta_k\}$ and $\{\zeta_k\}$ choices on Performance; $\Delta = 500$ and $p = 10$.

4.3. Impact of step size choices (expt.3)

The theoretical results presented in Section 3 show that the convergence rates of LLN, SA, and SAM estimators are affected by the choice of $\{\alpha_k\}$, $\{\eta_k\}$, and $\{\zeta_k\}$, respectively. Fig. 5 provides a numerical verification of the same. The details are as follows. We chose $\Delta = 500$ and $p = 10$. Notice that the page change rate is again very high, whereas the crawling frequency is relatively very low value. We then use the LLN estimator with three different choices of $\{\alpha_k\}$; these choices are shown in Fig. 5(a) itself. The LLN estimator with $\alpha_k = k^{0.75}$ has the worst performance. This behavior matches the prediction made by Theorem 1. In Fig. 5(b), we again consider the same setup as above. However, this time we run the SA estimator with three different choices of $\{\eta_k\}$; the choices are given in the figure itself. We see that the performance for $\eta = 0.5$ is better than the other cases.

We now analyze the impact of varying $\{\eta_k\}$ and $\{\zeta_k\}$ on the performance of the SAM estimator. Let ζ_k be of form given in (10). Based on our conjecture below (3), pick $\eta_k = (k + 1)^{-\eta}$ and $\beta_k = (k + 1)^{-\beta}$ with $\beta \in (0, 1]$ and $\beta < \eta < 2\beta$. In Fig. 5(c), we fix $\eta = 0.8$ and vary β ; these choices are shown in the figure itself. The SAM estimator with $\beta = 0.4$ reaches the limit very quickly, however, it is very noisy and keeps fluctuating around actual change rate. The fluctuations reduce as the value of β increases; however, larger values of β also slow down the rate at which the error decreases. We observe that the SAM estimator with $\beta = 0.6$ has the best performance. In Fig. 5(d), we fix $\beta = 0.6$ and vary η . The figure seems to suggest that a larger η increases the convergence rate but, simultaneously, also increases the fluctuations.

4.4. Practical recommendations

Here, we provide some recommendations on which estimator to use in practice. Our conclusions are based on what we have observed in the numerical experiments discussed in Section 4. We summarize them as follows.

- *High frequency crawling*: If the crawling frequency p is comparable to Δ , all estimators (LLN, SA, SAM and MLE) perform well except the Naive estimator. However, we do not recommend MLE as it is offline and very time-consuming. The examples that correspond to this scenario are depicted in Figs. 1(b) and 2.
- *Low frequency crawling*: There are two sub-cases depending on the value of p as compared to Δ .
 - Relatively very low p : The Naive estimator is very bad for this scenario as there will be several missed changes which will be unaccounted for. We recommend LLN or SAM estimator as they both outperform SA estimator; the example that corresponds to this scenario is depicted in Fig. 3. For similar reasons as in the previous case, we do not recommend the MLE estimator.
 - Relatively moderate p : The Naive estimator is again a bad choice here. Amongst the rest, we recommend the LLN estimator when several l_k values are available. Otherwise, one can use SAM or the MLE estimator; the offline nature of the MLE will be of concern here as well. The examples that correspond to this scenario are depicted in Figs. 1(c) and 4.

5. Estimating optimal crawling rates

In this section, we discuss how our estimators can be used to identify the optimal crawling rates. Formally, we suppose that a search engine's local cache consists of N pages. Let p_i denote the rate at which page i is crawled. The goal then is to find the optimal crawling rates such that the overall freshness of the local cache, i.e.,

$$\lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T \left(\sum_{i=1}^N w_i \mathbf{1}\{\text{Fresh}(i, t)\} \right) dt \right], \quad (17)$$

is maximized subject to the constraint $\sum_{i=1}^N p_i \leq B$. Here, $T > 0$ is the time horizon, w_i denotes the importance of the i th page, $B \geq 0$ is a bound on the overall crawling frequency, $\mathbf{1}\{\text{Fresh}(i, t)\}$ is the indicator that page i is fresh at time t , i.e., the local copy matches the actual page.

In [11], it was shown that maximizing (17) under a bandwidth constraint for large enough T corresponds to maximizing $F(p) = \sum_{i=1}^N (w_i p_i / (p_i + \Delta_i))$, where $p \equiv (p_1, \dots, p_N)$. Importantly, it was shown there that this latter optimization problem can be solved efficiently (in $O(N \log N)$ iterations) and provided an algorithm for the same. However, that algorithm requires that the Δ_i 's be known in advance. Our goal here is to combine their algorithm with our estimators and try and determine the optimal crawling rates.

Taking inspiration from [20], we consider the following hypothetical setup. We consider $N = 50$ pages, in which we presume that there are $\sqrt{50} \approx 7$ pages that change very frequently, i.e., they account for (say) 90% of the total changes in the system. Accordingly, we suppose that the change rate Δ_i for each frequently changing page is $4.5/7$, while for the others it is $0.5/43$. We further assume that the bound on the overall bandwidth is $B = 5$. We further assume that the frequently changing pages are more important and assign uniform weight of 2. On the other hand every other page is presumed to have uniform weight of 1.

We then use the following strategy. We arbitrarily initialize $p_i = B/N = 0.1$ for all i , i.e., B is uniformly divided across all the N pages. Since these p_i values are arbitrarily chosen, these need not be the optimal crawling rates. Thereafter, we run each of our estimators for (say) 50 iterations. We then use the estimates of Δ_i at the 50th iteration as input to [11, Algorithm 2] and obtain the associated possibly sub-optimal crawling rates. Denoting these new rates as p_i again, we now repeat the above procedure. That is, we use p_i 's for 50 iterations to estimate the Δ_i 's and, in turn, use the later to obtain estimates for the new p_i 's.

Fig. 6 compares the estimated crawling rates obtained using our three estimators with the optimal ones obtained by using the actual Δ_i values in [11, Algorithm 2] for the two kinds of pages. In this experiment, the two SA-based estimators appear to perform better than the LLN estimator, at least initially. Note that p_5 denotes the optimal crawling rate for the 5-th page, while p_5^{SA} denotes the estimate obtained by using the SA estimator, etc. p_{10} , p_{10}^{SA} , etc. have similar meanings in relation to the 10-th page. The parameters we chose for our different estimators are as follows. For the LLN estimator, we chose $\alpha_k \equiv 1$. For the SA estimator, we chose $\eta_k = (k + 1)^{-\eta}$ with $\eta = 0.75$. For the SAM estimator, we chose $\eta_k = (k + 1)^{-\eta}$ with $\eta = 1.3$ and $\zeta_k = (\beta_k - \omega \eta_k) / \beta_{k-1}$ with $\omega = 1$, and $\beta_k = (k + 1)^{-\beta}$ for $\beta = 0.75$.

6. Conclusion and future work

We propose three new online approaches for estimating the rate of change of web pages. We provide theoretical guarantees for their convergence and also provide numerical simulations to compare their performances. From experiments, one can verify that the proposed estimators perform significantly better than the Naive estimator. Also, they have

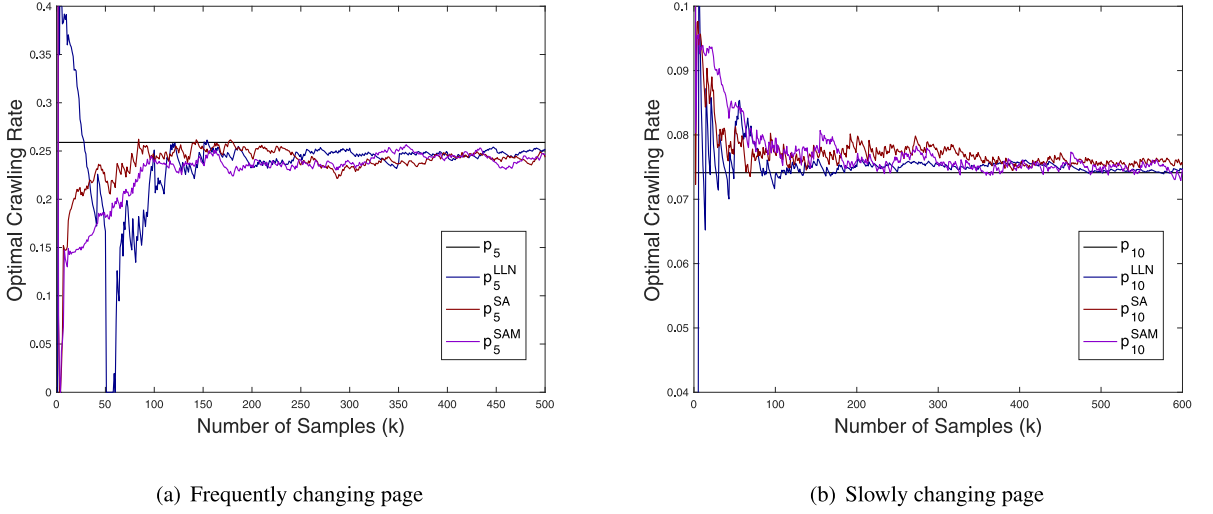


Fig. 6. Adaptive estimation of the optimal crawling rate.

extremely simple update rules which make them computationally attractive when compared to MLE. We also provide important insights on which estimator one should use in practice.

The performance of both our estimators currently depend on the choice of $\{\alpha_k\}$, $\{\eta_k\}$, and $\{\zeta_k\}$ respectively. One aspect to analyze in the future would be to ask what would be the ideal choice for these sequences that would help attain the fastest convergence rate. Another interesting research direction to pursue is to combine the online estimation with dynamic optimization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work is partly supported by ANSWER project PIA FSN2 (P15 9564-266178 \DOS0060094) and DST-Inria project “Machine Learning for Network Analytics” IFC/DST-Inria-2016-01/448. Research of Gugan Thoppe is supported by IISc’s, India start up grants SG/MHRD-19-0054 and SR/MHRD-19-0040. The authors would also like to thank A. Budhiraja for several useful discussions concerning [Theorem 3](#).

Appendix. Convergence of stochastic approximation algorithms

In this section, we discuss results from literature that provide sufficient conditions for convergence of both one-timescale and two-timescale stochastic approximation algorithms.

We begin by discussing the convergence of a generic one-timescale stochastic approximation algorithm. This result is obtained by combining [\[41, Chapter 2, Corollary 4,\]](#) and [\[41, Chapter 3, Theorem 7\]](#).

Theorem 9 (Convergence of One-timescale Stochastic Approximation [\[41\]](#)). Consider the update rule

$$y_{k+1} = y_k + \eta_k[h(y_k) + M_{k+1}],$$

where η_k is a positive scalar; $y_k, M_k \in \mathbb{R}^d$; and $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a deterministic function. Suppose the following conditions hold:

- (i.) $\sum_{k=0}^{\infty} \eta_k = \infty$ and $\sum_{k=0}^{\infty} \eta_k^2 < \infty$.
- (ii.) $\{M_k\}$ is a martingale difference sequence with respect to the increasing family of σ -fields

$$\mathcal{F}_k := \sigma(y_j, M_j, j \leq k), \quad k \geq 0.$$

That is, $\mathbb{E}[M_{k+1} | \mathcal{F}_k] = 0$ a.s., $k \geq 0$. Further, there is a constant $C \geq 0$ such that $\mathbb{E}[\|M_{k+1}\|^2 | \mathcal{F}_k] \leq C(1 + \|y_k\|^2)$ a.s. for all $k \geq 0$.

- (iii.) h is a globally Lipschitz continuous function. Further, the ODE $\dot{y}(t) = h(y(t))$ has a unique globally asymptotically stable equilibrium y_* .

(iv.) There exists a continuous function $h_\infty : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that the functions $h_c(x) := h(cx)/c$, $c \geq 1$, satisfy $h_c \rightarrow h_\infty$ uniformly on compact sets as $c \rightarrow \infty$. Further, the ODE $\dot{y}(t) = h_\infty(y(t))$ has the origin as its unique globally asymptotically stable equilibrium.

Then, $y_k \rightarrow y_*$ a.s.

Oftentimes, stochastic approximation algorithms contain an additional perturbation term that is asymptotically negligible. The next result discusses convergence of such algorithms.

Proposition 10 (Convergence of Perturbed One-timescale Stochastic Approximation). Consider the update rule

$$y_{k+1} = y_k + \eta_k[h(y_k) + \epsilon_k + M_{k+1}],$$

where ϵ_k is an additional perturbation term while the other terms have the same meaning as in Theorem 9. Suppose that the four conditions listed in Theorem 9 hold true. Further, suppose $\|\epsilon_k\| \leq C\rho_k(1 + \|y_k\|)$ a.s. for $k \geq 0$, where C is a positive constant and $\{\rho_k\}$ is a sequence of positive scalars such that $\lim_{k \rightarrow \infty} \rho_k = 0$. Then, $y_k \rightarrow y_*$ a.s.

Proof. We only give a sketch of the proof since the arguments are more or less similar to the ones used to derive Theorem 9. As mentioned before, this latter result follows from [41, Chapter 2, Corollary 4] and [41, Chapter 3, Theorem 7]. We now briefly discuss how, even in the presence of the additional perturbation term, these two results continue to hold.

- [41, Chapter 2, Corollary 4]: This result follows from [41, Chapter 2, Theorem 2] which, in turn, follows from [41, Chapter 2, Lemma 1]. However, as shown in extension 3 in [41, pg. 17], this latter result goes through even in the presence of the perturbation term $\{\epsilon_k\}$. This is because ϵ_k is asymptotically negligible a.s. More specifically, observe that the sequence $\{y_k\}$ is a.s. bounded under assumption (A4) given on [41, pg. 17]. This implies that $\{\epsilon_k\}$ is a random bounded sequence which is $o(1)$ a.s.; the latter is true since $\rho_k \rightarrow 0$.
- [41, Chapter 3, Theorem 7]: The proof of this result is based on Lemmas 1 to 6 in [41, Chapter 3]. The first three of these lemmas concern the behavior of the solution trajectories of the limiting ODE $\dot{y}(t) = h_\infty(y(t))$. Since the perturbation term does not affect the definition of this limiting ODE in any way whatsoever, these three results continue to hold as before. Similarly, Lemma 5 in *ibid* is unchanged since it only concerns the convergence of the sum of martingale differences $\sum_k \eta_k \hat{M}_{k+1}$ (recall that the stepsize sequence in our update rule is η_k). With regard to the proof of Lemma 4 in *ibid*, observe that our update rule satisfies

$$\hat{y}(t(k+1)) = \hat{y}(t(k)) + \eta_k(h_{r(n)}(\hat{y}(t(k))) + \hat{\epsilon}_k + \hat{M}_{k+1}), \quad m(n) \leq k \leq m(n+1),$$

where $\hat{\epsilon}_k = \epsilon_k/r(n)$ while the other notations are analogous to the ones defined in [41, Chapter 3]. Because $\|\epsilon_k\| \leq C\rho_k(1 + \|y_k\|)$, $\rho_k \rightarrow 0$, and $r(n) \geq 1$, it follows that

$$\|\hat{\epsilon}_k\| \leq C_1(1 + \|\hat{y}(t(k))\|)^2$$

for some positive constant C_1 . Note that this is in similar spirit to (3.2.5) in *ibid*. It is then easy to see that the rest of the proof goes through as before. This shows that [41, Chapter 3, Lemma 4] continues to be true even in the presence of the perturbation term. Using exactly the same bound for $\|\hat{\epsilon}_k\|$ obtained above, one can see that the arguments in the proof of Lemma 6 in *ibid* hold as well. Thus, [41, Chapter 3, Theorem 7] continues to hold, which is exactly what we wanted to establish.

The desired result now follows. \square

We next state a result that discusses the convergence of a generic two-timescale stochastic approximation algorithm. The proof of this result is based on [41, Chapter 6, Theorem 2] and [42, Theorem 10].

Theorem 11 (Convergence of Two-timescale Stochastic Approximation [41,42]). Consider the update rules

$$\begin{aligned} u_{k+1} &= u_k + \gamma_k[h(u_k, z_k) + M_{k+1}^{(1)}], \\ z_{k+1} &= z_k + \beta_k[g(u_k, z_k) + M_{k+1}^{(2)}], \end{aligned}$$

where γ_k and β_k are positive scalars; $u_k, z_k, M_k^{(1)}, M_k^{(2)} \in \mathbb{R}^d$; and $h, g : \mathbb{R}^{2d} \rightarrow \mathbb{R}^d$ are two deterministic functions. Suppose the following conditions hold:

- (i.) $\sum_{k \geq 0} \gamma_k = \sum_{k \geq 0} \beta_k = \infty$, $\sum_{k \geq 0} (\gamma_k^2 + \beta_k^2) < \infty$, and $\lim_{k \rightarrow \infty} \frac{\beta_k}{\gamma_k} = 0$.
- (ii.) $\{M_k^{(1)}\}$ and $\{M_k^{(2)}\}$ are martingale difference sequences with respect to the increasing σ -fields

$$\mathcal{F}_k := \sigma(u_j, z_j, M_j^{(1)}, M_j^{(2)}, j \leq k), \quad k \geq 0.$$

Further, there exists a constant $C \geq 0$ such that $\mathbb{E}[\|M_{k+1}^{(i)}\|^2 | \mathcal{F}_k] \leq C(1 + \|u_k\|^2 + \|z_k\|^2)$ for $i = 1, 2$ and $k \geq 0$.

- (iii.) h and g are globally Lipschitz continuous functions. For each fixed z , the ODE $\dot{u}(t) = h(u(t), z)$ has a unique globally asymptotically stable equilibrium $\phi(z)$, where $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is Lipschitz continuous. Further, the ODE $\dot{z}(t) = g(\phi(z(t)), z(t))$ has a unique globally asymptotically stable equilibrium z_* .
- (iv.) The functions $h_c(u, z) := h(cu, cz)/c$, $c \geq 1$, satisfy $h_c \rightarrow h_\infty$ as $c \rightarrow \infty$, uniformly on compacts for h_∞ . Also, for each fixed $z \in \mathbb{R}^d$, the limiting ODE $\dot{u}(t) = h_\infty(u(t), z)$ has a unique globally asymptotically stable equilibrium $\phi_\infty(z)$, where $\phi_\infty : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a Lipschitz map. Further, $\phi_\infty(0) = 0$. Separately, the functions $g_c(z) := g(\phi_\infty(z), cz)/c$, $c \geq 1$, satisfy $g_c \rightarrow g_\infty$ as $c \rightarrow \infty$, uniformly on compacts for some g_∞ . Also, the limiting ODE $\dot{z}(t) = g_\infty(z(t))$ has the origin as its unique globally asymptotically stable equilibrium.

Then, $(u_k, z_k) \rightarrow (\phi(z_*), z_*)$ a.s.

The last and final result of this section concerns the convergence of two-timescale stochastic approximation with perturbation terms that are asymptotically negligible.

Proposition 12 (Convergence of Perturbed Two-timescale Stochastic Approximation). Consider the update rules

$$\begin{aligned} u_{k+1} &= u_k + \gamma_k [h(u_k, z_k) + \epsilon_k^{(1)} + M_{k+1}^{(1)}] \\ z_{k+1} &= z_k + \beta_k [g(u_k, z_k) + \epsilon_k^{(2)} + M_{k+1}^{(2)}], \end{aligned}$$

where $\epsilon_k^{(1)}, \epsilon_k^{(2)}$ are additional perturbation terms while the other terms have the same meaning as in Theorem 11. Suppose that the four conditions listed in Theorem 11 hold true. Further, suppose $\|\epsilon_k^{(i)}\| \leq C\rho_k^{(i)}(1 + \|u_k\| + \|z_k\|)$ a.s. for $k \geq 0$ and $i = 1, 2$, where C is a positive constant and $\{\rho_k^{(i)}\}$, $i = 1, 2$, are sequences of positive scalars such that $\lim_{k \rightarrow \infty} \rho_k^{(i)} = 0$. Then, $(u_k, z_k) \rightarrow (\phi(z_*), z_*)$ a.s.

Proof. As stated before, this result follows from [41, Chapter 6, Theorem 2] and [42, Theorem 10]. We now briefly discuss how these results continue to hold even in the presence of the perturbation terms $\epsilon_k^{(1)}$ and $\epsilon_k^{(2)}$.

- [41, Chapter 6, Theorem 2]: This result, as well as [41, Chapter 6, Lemma 1] on which it relies, are essentially proved by defining suitable one-timescale stochastic approximation algorithms and then using convergence results concerning the latter. In our situation, both these will have additional perturbation terms that are asymptotically negligible. Consequently, by arguing as in the third extension given in [41, pg. 27], it can be shown that the asymptotic behavior of these two algorithms remains unchanged even in the perturbed setup. Therefore, it follows that the conclusions of [41, Chapter 6, Theorem 2] continue to hold as before.
- [42, Theorem 10]: This result is based on Lemmas 2 to 7 and Lemma 9 as well as Theorems 6 and 7 in *ibid*. Lemmas 2 to 5 in *ibid* concern the limiting ODEs described in condition (iv.) of Theorem 11 above. The definitions of these ODEs do not depend on the presence or absence of the perturbation terms. Therefore, the aforementioned four lemmas continue to hold as before. On the other hand, Lemmas 6 and 9 in *ibid* rely on the results in Chapter 3 and Chapter 6 of [41]. As argued before, these results continue to hold even in the presence of perturbation terms and, consequently, so do Lemmas 6 and 9 in *ibid*. Finally, Theorems 8 and 10 in *ibid* build upon these seven Lemmas. Therefore, they hold as well in the perturbed setup.

The desired result now follows. \square

References

- [1] A. Heydon, M. Najork, Mercator: A scalable, extensible web crawler, *World Wide Web* 2 (4) (1999) 219–229.
- [2] C. Castillo, Effective web crawling, in: *Acm Sigir Forum*, vol. 39, Acm New York, NY, USA, Association for Computing Machinery, New York, NY, USA, 2005, pp. 55–56.
- [3] C. Olston, M. Najork, et al., Web crawling, *Found. Trends[®] Inf. Retr.* 4 (3) (2010) 175–246.
- [4] J. Edwards, K. McCurley, J. Tomlin, An adaptive model for optimizing performance of an incremental web crawler, in: *Proceedings of the 10th International Conference on World Wide Web*, vol.8, Association for Computing Machinery, New York, NY, USA, 2001, pp. 106–113.
- [5] K. Avrachenkov, A. Dudin, V. Klimenok, P. Nain, O. Semenova, Optimal threshold control by the robots of web search engines with obsolescence of documents, *Comput. Netw.* 55 (8) (2011) 1880–1893.
- [6] J. Cho, H. Garcia-Molina, The evolution of the web and implications for an incremental crawler, in: *26th International Conference on Very Large Databases*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000, pp. 1–18.
- [7] J. Cho, H. Garcia-Molina, Synchronizing a database to improve freshness, *ACM Sigmod Rec.* 29 (2) (2000) 117–128.
- [8] N. Matloff, Estimation of internet file-access/modification rates from indirect data, *ACM Trans. Model. Comput. Simul.* 15 (3) (2005) 233–253.
- [9] B.E. Brewington, G. Cybenko, How dynamic is the web? *Comput. Netw.* 33 (1–6) (2000) 257–276.
- [10] J. Cho, H. Garcia-Molina, Effective page refresh policies for web crawlers, *ACM Trans. Database Syst.* 28 (4) (2003) 390–426.
- [11] Y. Azar, E. Horvitz, E. Lubetzky, Y. Peres, D. Shahaf, Tractable near-optimal policies for crawling, *Proc. Natl. Acad. Sci.* 115 (32) (2018) 8099–8103.
- [12] A. Kolobov, Y. Peres, E. Lubetzky, E. Horvitz, Optimal freshness crawl under politeness constraints, in: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019, pp. 495–504.
- [13] K.E. Avrachenkov, V.S. Borkar, Whittle index policy for crawling ephemeral content, *IEEE Trans. Control Netw. Syst.* 5 (1) (2016) 446–455.
- [14] J. Niño Mora, A dynamic page-refresh index policy for web crawlers, in: *Analytical and Stochastic Modeling Techniques and Applications*, Springer International Publishing, Cham, 2014, pp. 46–60.
- [15] F. Douglass, A. Feldmann, B. Krishnamurthy, J.C. Mogul, Rate of change and other metrics: A live study of the World Wide Web, in: *USENIX Symposium on Internet Technologies and Systems*, vol. 119, 1997, pp. 35.

- [16] C.E. Wills, M. Mikhailov, Towards a better understanding of web resources and server responses for improved caching, *Comput. Netw.* 31 (11–16) (1999) 1231–1243.
- [17] A. Wolman, M. Voelker, N. Sharma, N. Cardwell, A. Karlin, H.M. Levy, On the scale and performance of cooperative web proxy caching, in: *Proceedings of the Seventeenth ACM Symposium on Operating Systems Principles*, 1999, pp. 16–31.
- [18] J. Cho, H. Garcia-Molina, Estimating frequency of change, *ACM Trans. Internet Technol.* 3 (3) (2003) 256–290.
- [19] U. Upadhyay, R. Busa-Fekete, W. Kotlowski, D. Pal, B. Szorenyi, Learning to crawl, in: *Thirty-Fourth AAAI Conference on Artificial Intelligence*, AAAI Press, New York, NY, USA, 2020, pp. 8471–8478.
- [20] M. Andrews, S. Borst, J. Lee, E. Martin-Lopez, K. Palyutina, Tracking the state of large dynamic networks via reinforcement learning, in: *IEEE INFOCOM 2020—IEEE Conference on Computer Communications*, IEEE, 2020, pp. 416–425.
- [21] S. Gadat, F. Panloup, S. Saadane, Stochastic heavy ball, *Electron. J. Stat.* 12 (1) (2018) 461–529.
- [22] B.T. Polyak, Some methods of speeding up the convergence of iteration methods, *Ussr Comput. Math. Math. Phys.* 4 (5) (1964) 1–17.
- [23] K. Avrachenkov, K. Patil, G. Thoppe, Change rate estimation and optimal freshness in web page crawling, in: *Proceedings of the 13th EAI International Conference on Performance Evaluation Methodologies and Tools*, 2020, pp. 3–10.
- [24] B.E. Brewington, G. Cybenko, Keeping up with the changing web, *Computer* 33 (5) (2000) 52–58.
- [25] S.R. Singh, Estimating the rate of web page updates, in: *Proc. International Joint Conferences on Artificial Intelligence*, ACM, San Francisco, CA, USA, 2007, pp. 2874–2879.
- [26] A. Kolobov, Y. Peres, C. Lu, E.J. Horvitz, Staying up to date with online content changes using reinforcement learning for scheduling, in: *Advances in Neural Information Processing Systems*, 2019, pp. 581–591.
- [27] G. Dalal, G. Thoppe, B. Szörényi, S. Mannor, Finite sample analysis of two-timescale stochastic approximation with applications to reinforcement learning, in: *Conference on Learning Theory*, PMLR, 2018, pp. 1199–1233.
- [28] G. Dalal, B. Szörényi, G. Thoppe, A tale of two-timescale reinforcement learning with the tightest finite-time bound, in: *Thirty-Fourth AAAI Conference on Artificial Intelligence*, AAAI Press, San Francisco, CA, USA, 2020, pp. 3701–3708.
- [29] M. Kaledin, E. Moulines, A. Naumov, V. Tadic, H.-T. Wai, Finite time analysis of linear two-timescale stochastic approximation with Markovian noise, 2020, arXiv preprint [arXiv:2002.01268](https://arxiv.org/abs/2002.01268).
- [30] G. Dalal, B. Szörényi, G. Thoppe, S. Mannor, Finite sample analyses for td (0) with function approximation, in: *Thirty-Second AAAI Conference on Artificial Intelligence*, AAAI Press, San Francisco, CA, USA, 2018, pp. 6144–6160.
- [31] I. Gitman, H. Lang, P. Zhang, L. Xiao, Understanding the role of momentum in stochastic gradient methods, 2019, arXiv preprint [arXiv:1910.13962](https://arxiv.org/abs/1910.13962).
- [32] T. Yang, Q. Lin, Z. Li, Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization, 2016, arXiv preprint [arXiv:1604.03257](https://arxiv.org/abs/1604.03257).
- [33] N.S. Aybat, A. Fallah, M. Gurbuzbalaban, A. Ozdaglar, Robust accelerated gradient methods for smooth strongly convex functions, *SIAM J. Optim.* 30 (1) (2020) 717–751.
- [34] M. Laborde, A. Oberman, A Lyapunov analysis for accelerated gradient methods: From deterministic to stochastic case, in: *International Conference on Artificial Intelligence and Statistics*, PMLR, 2020, pp. 602–612.
- [35] M. Assran, M. Rabbat, On the convergence of Nesterov's accelerated gradient method in Stochastic settings, 2020, arXiv preprint [arXiv:2002.12414](https://arxiv.org/abs/2002.12414).
- [36] A. Kulunchakov, J. Mairal, Estimate sequences for variance-reduced stochastic composite optimization, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 3541–3550.
- [37] B. Can, M. Gurbuzbalaban, L. Zhu, Accelerated linear convergence of stochastic momentum methods in wasserstein distances, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 891–901.
- [38] M. Cohen, J. Diakonikolas, L. Orecchia, On acceleration with noise-corrupted gradients, in: *International Conference on Machine Learning*, PMLR, 2018, pp. 1019–1028.
- [39] S. Vaswani, F. Bach, M. Schmidt, Fast and faster convergence of sgd for over-parameterized models and an accelerated perceptron, in: *The 22nd International Conference on Artificial Intelligence and Statistics*, PMLR, 2019, pp. 1195–1204.
- [40] N. Loizou, P. Richtárik, Momentum and stochastic momentum for stochastic gradient, newton, proximal point and subspace descent methods, *Comput. Optim. Appl.* 77 (3) (2020) 653–710.
- [41] V.S. Borkar, *Stochastic Approximation: A Dynamical Systems Viewpoint*, vol. 48, Springer, India, 2009.
- [42] C. Lakshminarayanan, S. Bhatnagar, A stability criterion for two timescale stochastic approximation schemes, *Automatica* 79 (2017) 108–114.