

General Temporally-Biased Sampling Schemes for Online Model Management

BRIAN HENTSCHEL, Harvard University, USA

PETER J. HAAS, University of Massachusetts Amherst, USA

YUANYUAN TIAN, IBM Research – Almaden, USA

To maintain the accuracy of supervised learning models in the presence of evolving data streams, we provide temporally-biased sampling schemes that weight recent data most heavily, with inclusion probabilities for a given data item decaying over time according to a specified “decay function”. We then periodically retrain the models on the current sample. This approach speeds up the training process relative to training on all of the data. Moreover, time-biasing lets the models adapt to recent changes in the data while—unlike in a sliding-window approach—still keeping some old data to ensure robustness in the face of temporary fluctuations and periodicities in the data values. In addition, the sampling-based approach allows existing analytic algorithms for static data to be applied to dynamic streaming data essentially without change. We provide and analyze both a simple sampling scheme (T-TBS) that probabilistically maintains a target sample size and a novel reservoir-based scheme (R-TBS) that is the first to provide both control over the decay rate and a guaranteed upper bound on the sample size. If the decay function is exponential, then control over the decay rate is complete, and R-TBS maximizes both expected sample size and sample-size stability. For general decay functions, the actual item inclusion probabilities can be made arbitrarily close to the nominal probabilities, and we provide a scheme that allows a trade-off between sample footprint and sample-size stability. R-TBS rests on the notion of a “fractional sample” and allows for data arrival rates that are unknown and time varying (unlike T-TBS). The R-TBS and T-TBS schemes are of independent interest, extending the known set of unequal-probability sampling schemes. We discuss distributed implementation strategies; experiments in Spark illuminate the performance and scalability of the algorithms, and show that our approach can increase machine learning robustness in the face of evolving data.

CCS Concepts: • **Information systems** → **Stream management**; **Data streaming**; • **Theory of computation** → *Distributed algorithms*; Sketching and sampling; • **Computing methodologies** → *Online learning settings*; *Distributed algorithms*;

Additional Key Words and Phrases: temporally-biased sampling, time-biased sampling, temporal decay, reservoir sampling, online model management, model robustness, fractional sample, latent sample, unequal-probability sampling, distributed sampling, model retraining

ACM Reference Format:

Brian Hentschel, Peter J. Haas, and Yuanyuan Tian. 2018. General Temporally-Biased Sampling Schemes for Online Model Management. *ACM Trans. Datab. Syst.* 1, 1, Article 1 (January 2018), 45 pages. <https://doi.org/0000001.0000001>

Authors’ addresses: Brian Hentschel, Harvard University, Cambridge, Massachusetts, USA, bhentschel@g.harvard.edu; Peter J. Haas, University of Massachusetts Amherst, Amherst, Massachusetts, USA, phaas@cs.umass.edu; Yuanyuan Tian, IBM Research – Almaden, San Jose, California, USA, ytian@us.ibm.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0362-5915/2018/1-ART1 \$15.00

<https://doi.org/0000001.0000001>

1 INTRODUCTION

A key challenge for machine learning (ML) is to keep ML models from becoming stale in the presence of evolving data. In the context of the emerging Internet of Things (IoT), for example, the data comprises dynamically changing sensor streams [34], and a failure to adapt to changing data can lead to a loss of predictive power.

One way to deal with this problem is to re-engineer existing static supervised learning algorithms to become adaptive. Some parametric methods—such as support-vector machines (SVMs) without the use of kernel methods, hidden Markov models, and regression models—can indeed be re-engineered so that the parameters are time-varying, but for many popular non-parametric algorithms such as k -nearest neighbors (kNN) classifiers, decision trees, random forests, gradient boosted machines, and so on, it is not at all clear how re-engineering can be accomplished. The 2017 Kaggle Data Science Survey [33] indicates that a substantial portion of the models that developers use in industry are non-parametric. We therefore consider alternative approaches in which we periodically retrain ML models, allowing static ML algorithms to be used in dynamic settings essentially as-is. There are several possible retraining approaches.

Retraining on cumulative data: Periodically retraining a model on all of the data that has arrived so far is clearly infeasible because of the huge volume of data involved. Moreover, recent data is swamped by the massive amount of past data, so the retrained model is not sufficiently adaptive.

Sliding windows: A simple sliding-window approach would be to, e.g., periodically retrain on the data from the last two hours. If the data arrival rate is high and there is no bound on memory, then one must deal with long retraining times caused by large amounts of data in the window. The simplest way to bound the window size is to retain the last n items. Alternatively, one could try to subsample within the time-based window [16]. The fundamental problem with all of these bounding approaches is that old data is completely forgotten; the problem is especially severe when the data arrival rate is high. This can undermine the robustness of an ML model in situations where old patterns can reassert themselves. For example, a singular event such as a holiday, stock market drop, or terrorist attack can temporarily disrupt normal data patterns, which will reestablish themselves once the effect of the event dies down. Periodic data patterns can lead to the same phenomenon. Another example, from [35], concerns influencers on Twitter: a prolific tweeter might temporarily stop tweeting due to travel, illness, or some other reason, and hence be completely forgotten in a sliding-window approach. Indeed, in real-world Twitter data, almost a quarter of top influencers were of this type, and were missed by a sliding window approach.

Temporally-biased sampling: An appealing alternative is a temporally biased sampling-based approach, i.e., maintaining a sample that heavily emphasizes recent data but also contains a small amount of older data, and periodically retraining a model on the sample. By using a time-biased sample, the retraining costs can be held to an acceptable level while not sacrificing robustness in the presence of recurrent patterns. This approach was proposed in [35] in the setting of graph analysis algorithms, and has recently been adopted in the MacroBase system [3]. The orthogonal problem of choosing when to retrain a model is also an important question, and is related to, e.g., the literature on “concept drift” [15]; in this paper we focus on the problem of how to efficiently maintain a time-biased sample.

In more detail, our time-biased sampling algorithms ensure that the “appearance probability” for a given data item—i.e., the probability that the item appears in the current sample—decays over time at a controlled rate. Specifically, we assume that items arrive in *batches* $\mathcal{B}_1, \mathcal{B}_2, \dots$, at time points t_1, t_2, \dots , where each batch contains 0 or more items and $t_k \rightarrow \infty$ as $k \rightarrow \infty$. Our goal is to generate a sequence $\{S_k\}_{k \geq 1}$, where S_k is a sample of the items that have arrived at or prior to

time t_k , i.e., a sample of the items in $\mathcal{U}_k = \bigcup_{i=1}^k \mathcal{B}_i$. These samples should be biased towards recent items, in the following sense. For $1 \leq i \leq k$, denote by $\alpha_{i,k} = t_k - t_i$ the *age* at time t_k of an item belonging to batch \mathcal{B}_i . Then for arbitrary times $t_i \leq t_j$ and items $x \in \mathcal{B}_i$ and $y \in \mathcal{B}_j$,

$$\Pr[x \in S_k] / \Pr[y \in S_k] = f(\alpha_{i,k}) / f(\alpha_{j,k}), \quad (1)$$

for any batch arrival time $t_k \geq t_j$, where f is a nonnegative and nonincreasing *decay function* such that $f(0) = 1$. Thus items with a given timestamp are sampled uniformly, and items with different timestamps are handled in a carefully controlled manner, such that the appearance probability for an item of age α is proportional to $f(\alpha)$. The criterion in (1), which is expressed in terms of wall-clock time, is natural and appealing in applications and, importantly, is interpretable and understandable to users.

Choosing a decay function: Although our primary focus is on developing sampling methods that can support a variety of decay functions, the question of how to choose a good decay function f is important, and a topic of ongoing research. Cohen and Strauss [9] discuss the choice of decay functions in the setting of time-decayed aggregates in telecommunications networks, and argue that the proper choice of a decay function depends on domain knowledge. In an example involving a reliability comparison between two telecommunications links, they conclude that a polynomial decay function best matches their intuition on how the comparison should evolve over time. More generally, the authors assert that there is a trade-off between the ability to decay quickly in the short term and the ability to potentially retain older data, and that the choice should depend on the perceived importance of older data and on the time scales of correlations between values; e.g., the latter might correspond to the amount of time it takes a prior pattern to reassert itself. The authors therefore argue for supporting a rich class of decay functions. Similarly, Xie et al. [35] show how a decay function can be chosen to meet application-specific criteria. For example, by using an exponential decay function $f(\alpha) = \exp(-\lambda\alpha)$ with $\lambda = 0.058$, a data item from 40 batches ago is $1/10$ as likely to appear in the current analysis as a newly arrived item. If training data is available, λ can also be chosen to maximize accuracy of a specified ML model via cross validation combined with grid search—in our experiments, where ground truth data was available, we found empirically that accuracy tended to be a quasiconvex function of λ , which bodes well for automatic optimization methods such as stochastic gradient descent. We find exponential and sub-exponential decay functions such as polynomial decay to be of the greatest interest. As will become apparent, exponential decay functions, though of limited flexibility, are the easiest to work with, and most prior work has centered around exponential decay. Super-exponential decay functions are of less practical interest: older items decay too fast and the sampling scheme behaves essentially like a sliding window.

Sample-size control: For the case in which the item-arrival rate is high, the main issue is to keep the sample size from becoming too large. On the other hand, when the incoming batches become very small or widely spaced, the sample sizes for all of the time-biased algorithms that we discuss (as well as for sliding-window schemes based on wall-clock time) can become small. This is a natural consequence of treating recent items as more important, and is characteristic of any sampling scheme that satisfies (1). We emphasize that—as shown in our experiments—a smaller, but carefully time-biased sample typically yields better prediction accuracy than a sample that is larger due to overloading with too much recent data or too much old data. I.e., more sample data is not always better. Indeed, with respect to model management, this decay property can be viewed as a feature in that, if the data stream dries up and the sample decays to a very small size, then this is a signal that there is not enough new data to reliably retrain the model, and that the current version should be kept for now. In any case, we show that, within the class of sampling

algorithms that support exponential decay, our new R-TBS algorithm with an exponential decay function maximizes the expected sample size whenever the sample is not saturated.

Prior work: It is surprisingly hard to both enforce (1) and to bound the sample size. As discussed in detail in Section 7, prior algorithms cannot handle arbitrary decay functions, and can only support “forward decay” schemes, which are less intuitive for users than “backward decay” schemes that enforce (1), and can lead to both numerical issues and poor adaptation behavior for ML algorithms. The only decay functions that support (1) and are handled by prior algorithms are the exponential decay functions, because backward and forward decay coincide in this case. Even in this restricted setting, prior algorithms that bound the sample size either cannot consistently enforce (1) or cannot handle wall-clock time. Examples of the former include algorithms based on the A-Res scheme of Efraimidis and Spirakis [14] and Chao’s algorithm [7]. A-Res enforces conditions on the *acceptance* probabilities of items; this leads to appearance probabilities that, unlike (1), are both hard to compute and not intuitive. In the Online Appendix we demonstrate how Chao’s algorithm can be specialized to the case of exponential decay and modified to handle batch arrivals. We then observe that the resulting algorithm fails to enforce (1) either when initially filling up an empty sample or in the presence of data that arrives slowly relative to the decay rate, and hence fails if the data rate fluctuates too much. The second type of algorithm, due to Aggarwal [1], can only control appearance probabilities based on the indices of the data items. For example, after n items arrive, one could require that, for some specified $k < n$, the $(n - k)$ th item is $1/10$ as likely to be in the sample as the current item. If the data arrival rate is constant, then this might correspond to a constraint of the form “a data item that arrived 10 hours ago is $1/10$ as likely to be in the sample as the current item”. For varying arrival rates, however, it is impossible to enforce the latter type of constraint, and a large batch of arriving data can prematurely flush out older data. Thus our new sampling schemes are interesting in their own right, significantly expanding the set of unequal-probability sampling techniques.

T-TBS: We first provide and analyze Targeted-Size Time-Biased Sampling (T-TBS), a relatively simple algorithm that generalizes the Bernoulli sampling scheme in [35]. T-TBS allows complete control over the decay rate (expressed in wall-clock time) and probabilistically maintains a target sample size. That is, the expected and average sample sizes converge to the target and the probability of large deviations from the target decreases exponentially or faster in both the target size and the deviation size. T-TBS is easy to implement and highly scalable when applicable, but only works under the strong restriction that the mean sizes of the arriving batches are constant over time and known a priori. T-TBS is a good choice in some scenarios (see Section 3), but many applications have non-constant, unknown mean batch sizes, thus cannot tolerate sample overflows.

R-TBS: We then provide a novel algorithm, Reservoir-Based Time-Biased Sampling (R-TBS), that is the first to simultaneously enforce (1) at all times, provide a guaranteed upper bound on the sample size, and allow unknown, varying data arrival rates. Guaranteed bounds are desirable because they avoid memory management issues associated with sample overflows, especially when large numbers of samples are being maintained—so that the probability of *some* sample overflowing is high—or when sampling is being performed in a limited memory setting such as at the “edge” of the IoT. Also, bounded samples reduce variability in retraining times and do not impose upper limits on the incoming data flow. For an exponential decay function, the appearance probability for an item of age α is always exactly proportional to $f(\alpha)$; in general, for a “cutoff age” α^* , the appearance probability for an item of age α is proportional to $\tilde{f}(\alpha)$, where $\tilde{f}(\alpha) = f(\alpha)$ for $\alpha \leq \alpha^*$ but $\tilde{f}(\alpha) \neq f(\alpha)$ for $\alpha > \alpha^*$. At the cost of additional storage, the user can make α^* arbitrarily large—so that only a small set of very old items are affected—and the discrepancy $|f - \tilde{f}|$ for these old items arbitrarily small. We emphasize that, even though R-TBS involves some

approximations in the case of general decay functions, the magnitude of these approximations is completely controllable a priori by the user; in contrast, prior schemes such as A-Res and Chao's Algorithm offer no control over departures from (1) and indeed it can be difficult even to quantify the extent of these departures.

The idea behind R-TBS is to adapt the classic reservoir sampling algorithm, which bounds the sample size but does not allow time biasing. Our approach rests on the notion of a "fractional" sample whose nonnegative size is real-valued in an appropriate sense. For exponential decay, we show that R-TBS maximizes the expected sample size whenever the data arrival rate is low and also minimizes the sample-size variability; in general, there again is a user-controllable tradeoff between storage requirements and sample size stability.

Distributed implementation: Both T-TBS and R-TBS can be parallelized. Whereas T-TBS is relatively straightforward to implement, an efficient distributed implementation of R-TBS is nontrivial. We exploit various implementation strategies to reduce I/O relative to other approaches, avoid unnecessary concurrency control, and make decentralized decisions about which items to insert into, or delete from, the reservoir.

Extensions of our prior work: A preliminary version of this work appeared in [18]; that paper focused entirely on the case of exponential decay and was missing many of the proofs for the given theoretical results. In the current paper, we extend our results to the setting of general decay functions. Handling such functions requires significant extensions to the algorithms, theory, and experimental study given in [18]. Interestingly, viewing the original R-TBS algorithm in [18] as a special case of the general algorithm has led to streamlining of the original algorithm as well as its theoretical analysis. The current paper (and online appendix) contains all relevant proofs.

Organization: The rest of the paper is organized as follows. In Section 2 we describe our batch-arrival model and, to provide context for the current work, discuss two prior simple sampling schemes: a simple Bernoulli scheme as in [35] and the classical reservoir sampling scheme, modified for batch arrivals. These methods either bound the sample size but do not control the decay rate, or control the decay rate but not the sample size. We next present and analyze the T-TBS and R-TBS algorithms in Section 3 and Section 4. We describe the distributed implementation in Section 5, and Section 6 contains experimental results. We review the related literature in Section 7 and conclude in Section 8.

2 BACKGROUND

For the remainder of the paper, we focus on settings in which batches arrive at regular time intervals, so that $t_i = i\Delta$ for some $\Delta > 0$. This simple integer batch sequence often arises from the discretization of time [30, 36]. Specifically, the continuous time domain is partitioned into intervals of length Δ , and the items are observed only at times $\{k\Delta : k = 1, 2, \dots\}$. All items that arrive in an interval $((k-1)\Delta, k\Delta]$ are treated as if they arrived at time $k\Delta$, i.e., at the end of the interval, so that all items in batch \mathcal{B}_i have time stamp $i\Delta$. It follows that the age at time t_k of an item that arrived at time $t_i \leq t_k$ is simply $\alpha_{i,k} = (k-i)\Delta$.

In this section, we briefly review two classical sampling schemes whose properties we will combine in the R-TBS algorithm.

Bernoulli Time-Biased Sampling (B-TBS): A well known, simple Bernoulli time-biased sampling scheme processes each incoming item, one at a time, by first downsampling the current sample and then accepting the incoming item into the sample with probability 1. Downsampling is accomplished by flipping a coin independently for each item in the sample: an item is retained in the sample with probability p and removed with probability $1-p$. To adapt this sampling scheme to our batch-arrival setting, we process incoming items a batch at a time, and implicitly assume an exponential decay function $f(\alpha) = e^{-\lambda\alpha}$, setting $p = e^{-\lambda\Delta}$. Moreover, we take advantage of the

ALGORITHM 1: Bernoulli time-biased sampling (B-TBS)

```

1  $\lambda$ : decay factor ( $\geq 0$ )
2 Initialize:  $S \leftarrow \emptyset$ ;  $p \leftarrow e^{-\lambda\Delta}$                                 //  $p$  = retention prob.
3 for  $i \leftarrow 1, 2, \dots$  do
4    $M \leftarrow \text{BINOMIAL}(|S|, p)$                                 // simulate  $|S|$  trials
5    $S \leftarrow \text{SAMPLE}(S, M)$                                 // retain  $M$  random elements
6    $S \leftarrow S \cup \mathcal{B}_i$ 
7   output  $S$ 

```

fact that the foregoing downsampling operation is probabilistically equivalent to pre-selecting the number M of items to retain according to a binomial distribution and then choosing the actual set of M retained items uniformly from the elements in the current sample; see the Online Appendix for a proof of this fact. Generating a sample of M can be done efficiently using standard algorithms [32], and obviates the need for executing multiple coin flips.

The resulting sampling scheme is given as Algorithm 1. At each time t_i we accept each incoming item $x \in \mathcal{B}_i$ into the sample with probability 1 (line 6). Downsampling is accomplished in lines 4 and 5: the function $\text{BINOMIAL}(j, r)$ returns a random sample from the binomial distribution with j independent trials and success probability r per trial, and the function $\text{SAMPLE}(A, m)$ returns a uniform random sample, without replacement, containing $\min(m, |A|)$ elements of the set A ; note that the function call $\text{SAMPLE}(A, 0)$ returns an empty sample for any empty or nonempty A .

To see that Algorithm 1 enforces the relation in (1) as required, observe that the sequence of samples is a set-valued Markov process, so that

$$\Pr[x \in S_{i+j} \mid x \in S_{i+j-1}, x \in S_{i+j-2}, \dots, x \in S_i] = \Pr[x \in S_{i+j} \mid x \in S_{i+j-1}]$$

for $i, j \geq 1$. We then have, for $x \in \mathcal{B}_i$,

$$\Pr[x \in S_k] = \Pr[x \in S_i] \times \prod_{j=1}^{k-i} \Pr[x \in S_{i+j} \mid x \in S_{i+j-1}] = 1 \times p^{k-i} = e^{-\lambda(k-i)\Delta} = e^{-\lambda(t_k - t_i)}, \quad (2)$$

and (1) follows immediately from (2). Thus Algorithm 1 precisely controls the relative inclusion probabilities according to the exponential decay function f given above. This is the algorithm used, e.g., in [35] to implement time-biased edge sampling in dynamic graphs.

Unfortunately, the user cannot independently control the expected sample size, which is completely determined by λ and the sizes of the incoming batches. In particular, if the batch sizes systematically grow over time, then sample size will grow without bound. Arguments in [35] show that if $\sup_i |\mathcal{B}_i| < \infty$, then the sample size can be bounded, but only probabilistically. Because B-TBS is a special case of T-TBS with an exponential decay function and a unitary acceptance probability for arriving items, the results in Section 3 represent a significant extension and refinement of the analysis in [35].

Batched Reservoir Sampling (B-RS): The classical reservoir sampling algorithm [22, 25] maintains a bounded uniform sample of items in a data stream. The idea is to fill up the reservoir with the first n items, where n is the reservoir size. For $k > n$, the k th incoming item is accepted into the sample with probability $q_k = n/k$, and an accepted item overwrites a randomly chosen victim. Our choice of q_k is intuitively motivated by the observation that, in general, a given item from a population of size k appears in a uniform sample of size $n \leq k$ with probability precisely equal to n/k .

ALGORITHM 2: Batched reservoir sampling (B-RS)

```

1  $n$ : maximum sample size
2 Initialize:  $S \leftarrow \emptyset$ ;  $W \leftarrow 0$ ; for  $i = 1, 2, \dots$  do
3    $C = \min(n, W + |\mathcal{B}_i|)$  //new sample size
4    $M \leftarrow \text{HYPERGEO}(C, |\mathcal{B}_i|, W)$ 
   //add  $M$  elements to  $S$ , overwriting  $\max(M - (n - |S|), 0)$  items
5    $S \leftarrow \text{SAMPLE}(S, \min(n - M, |S|)) \cup \text{SAMPLE}(\mathcal{B}_i, M)$ 
6    $W \leftarrow W + |\mathcal{B}_i|$ 
7   output  $S$ 

```

We can extend the classical algorithm to our batch setting; to our knowledge, a batch-oriented variant has not appeared previously in the literature. To informally motivate the algorithm, we generalize the foregoing intuition. For $k \geq 1$, let $\mathcal{U}_k = \bigcup_{j=1}^k \mathcal{B}_j$ be the set of items arriving up through time t_k and set $W_k = |\mathcal{U}_k|$. Suppose that the sample is full (i.e., $|S_{k-1}| = n$) just before batch \mathcal{B}_k arrives. After processing \mathcal{B}_k , we ought to have a uniform sample of n items from the set $\mathcal{U}_k = \mathcal{U}_{k-1} \cup \mathcal{B}_k$. We would thus expect the number M of \mathcal{B}_k -items in the sample to follow a hypergeometric($n, |\mathcal{B}_k|, W_{k-1}$) distribution; here the hypergeometric(k, a, b) probability mass function is given by $p(n) = \binom{a}{n} \binom{b}{k-n} / \binom{a+b}{k}$ if $\max(0, k-b) \leq n \leq \min(a, k)$ and $p(n) = 0$ otherwise. This motivates us to accept new items from \mathcal{B}_k into the sample by first generating the number of items to accept as a hypergeometric variate M and then selecting M specific items for acceptance in a random and uniform manner. As with classical reservoir sampling, incoming items arriving before the sample fills up are accepted into the sample with probability 1 and do not overwrite random victims, whereas subsequent incoming items do overwrite random victims. In the corner case where $W_{k-1} < n$ and $W_{k-1} + |\mathcal{B}_k| \geq n$, so that an incoming batch would cause the sample to overflow if all items were accepted, we generate M as before, but $\min(n - W_{k-1}, M)$ of these elements are accepted into the sample without overwriting a random victim, and the remainder overwrite a random victim from \mathcal{U}_{k-1} .

The resulting sampling scheme is given as Algorithm 2. In the algorithm, `SAMPLE` is defined as before and `HYPERGEO`(k, a, b) returns a sample from the hypergeometric(k, a, b) distribution; see [32] for a discussion of efficient implementations of `HYPERGEO`. The Online Appendix contains a formal proof of correctness. Although B-RS guarantees an upper bound on the sample size, it does not support time biasing in that all items seen so far are equally likely to be in the sample. The R-TBS algorithm (Section 4) maintains a bounded reservoir as in B-RS while simultaneously allowing time-biased sampling as in B-TBS.

3 TARGETED-SIZE TBS

As a first step towards time-biased sampling with a controlled sample size, we provide the T-TBS scheme, which improves upon B-TBS by ensuring the inclusion property in (1) while providing probabilistic guarantees on the sample size. Throughout, we focus on the case where the batch sizes $\{|\mathcal{B}_k|\}_{k \geq 1}$ are independent and identically distributed (i.i.d.) with common mean $b < \infty$, and assume that the decay function f satisfies $\lim_{\alpha \rightarrow \infty} f(\alpha) = 0$.

3.1 The Algorithm

The key idea is to not just downsample to remove older items as in B-TBS, but to also downsample incoming batches at a rate q such that n becomes (asymptotically) the “equilibrium” sample size. Unlike B-TBS, we now want to have the retention probability of an item depend on its age. In

particular, if for $x \in \mathcal{B}_i$, we set the retention probability at time $t_k \geq t_i$ equal to

$$p_{i,k} = f(\alpha_{i,k})/f(\alpha_{i,k-1}), \quad (3)$$

then we have

$$\Pr[x \in S_k] = \Pr[x \in S_i] \times \prod_{j=1}^{k-i} \Pr[x \in S_{i+j} \mid x \in S_{i+j-1}] = q \times \prod_{j=1}^{k-i} \frac{f(\alpha_{i,i+j})}{f(\alpha_{i,i+j-1})} = qf(\alpha_{i,k}), \quad (4)$$

and (1) follows immediately from (4).

To choose q , we reason as follows. Suppose that the sample size equals the target value n and we are about to process batch \mathcal{B}_k . Prior to incrementing the ages and processing the arriving batch, the ages in the sample range from $\alpha_{0,k-1} = t_{k-1}$ down to $\alpha_{k-1,k-1} = 0$, with an expected number $n\phi_{i,k-1}$ of sample items belonging to batch \mathcal{B}_i , where $\phi_{i,k-1} = f(\alpha_{i,k-1})/\sum_{j=0}^{k-1} f(\alpha_{j,k-1})$. Thus the expected number of \mathcal{B}_i items removed prior to processing \mathcal{B}_k is $n\phi_{i,k-1}(1 - p_{i,k})$, where $p_{i,k}$ is defined in (3). Summing over all batches \mathcal{B}_i , we find, after some algebra, that the expected total number of removed items is $n\gamma_k$, where

$$\gamma_k = \sum_{i=0}^{k-1} \phi_{i,k-1}(1 - p_{i,k}) = 1 - \frac{\sum_{i=0}^{k-1} f(\alpha_{i,k})}{\sum_{i=0}^{k-1} f(\alpha_{i,k-1})} \quad (5)$$

for $k \geq 1$. On the other hand, the expected number of items entering the sample is $q_k b$ (where we initially allow q to depend on k). For n to be an equilibrium point, we equate the expected inflow and outflow and solve for q_k to obtain $q_k = n\gamma_k/b$. If

$$\lim_{k \rightarrow \infty} \gamma_k = \gamma \quad (6)$$

for some $\gamma \in (0, 1]$, then $\lim_{k \rightarrow \infty} q_k = n\gamma/b$. In light of (4), we see that by setting $q = n\gamma/b$, we ensure that (1) holds at all times, and that the sample size n is asymptotically an equilibrium point as k , the number of batches processed, becomes large. (See Section 3.2 for a formal statement and proof.) Note that, even if we always accept all items in an arriving batch (i.e., $q = 1$) but the resulting expected inflow b is less than the expected outflow $n\gamma$, the sample will consistently fall below n , and so we require that $b \geq n\gamma$.

Because we are assuming that $t_i = i\Delta$ for $i \geq 1$, we can easily derive a necessary and sufficient condition for (6) to hold. Writing $f_i = f(i\Delta)$, we have $\gamma_k = (f_0 - f_k)/F_{k-1}$, where $F_j = \sum_{i=0}^j f_i$ for $j \geq 0$. Thus, if $F_\infty = \sum_{i=0}^\infty f_i < \infty$, then (6) holds with $\gamma = f_0/F_\infty = 1/F_\infty$, so that $q = n/(bF_\infty)$. Necessity follows from the fact that $\lim_{k \rightarrow \infty} f_k = 0$ by assumption. For polynomial decay with $f(\alpha) = 1/(1 + \alpha)^s$, we have $\gamma = \zeta(s, 1/\Delta)/\Delta^s$, where ζ is the Hurwitz zeta function. If the decay function is exponential, i.e., $f(\alpha) = \exp(-\lambda\alpha)$, and we choose a time scale so that $\Delta = 1$, then a simple calculation shows that $\gamma_k \equiv 1 - \exp(-\lambda)$ and $p_{i,k} \equiv \exp(-\lambda)$ for $k \geq i \geq 1$, and we obtain the T-TBS algorithm for exponential decay as described in [18]. Here n is an equilibrium point for every value of k , and not merely in an asymptotic sense as $k \rightarrow \infty$. For this special case, we do not need to maintain the arrival timestamp for each item, and therefore do not need to partition the sample items based on arrival time. If we further assume that $q = 1$, then we obtain the B-TBS algorithm as a special case in which the equilibrium sample size is $b/\gamma = b/(1 - e^{-\lambda})$, which is completely determined by b and λ . For functions for which γ cannot be solved analytically, we can compute γ numerically.

The resulting sampling scheme is given as Algorithm 3; it precisely controls inclusion probabilities in accordance with (1) while constantly pushing the sample size toward the target value n . We represent a sample S as a collection of sets H_i , where H_i is the set of sample items that arrived at time t_i ; thus $H_i \subseteq \mathcal{B}_i$. The function GAMMA in line 4 computes the constant γ defined above.

ALGORITHM 3: Targeted-size TBS (T-TBS)

```

1  $f$ : decay function
2  $n$ : target sample size
3  $b$ : assumed mean batch size such that  $b \geq n\gamma$ 
4 Initialize:  $S \leftarrow \emptyset$ ;  $\gamma = \text{GAMMA}(f)$ ;  $q \leftarrow n\gamma/b$ 
5 for  $k \leftarrow 1, 2, \dots$  do
6   for  $H_i \in S$  do                                     //update current sample items
7      $p = f(\alpha_{i,k})/f(\alpha_{i,k-1})$                        //retention probability
8      $m \leftarrow \text{BINOMIAL}(|H_i|, p)$                      //simulate  $|H_i|$  trials
9     if  $m > 0$  then
10       $H_i \leftarrow \text{SAMPLE}(H_i, m)$                    //retain  $m$  random elements
11    else                                                 //Hi is now empty
12       $S \leftarrow S \setminus \{H_i\}$ 
13   $l \leftarrow \text{BINOMIAL}(|\mathcal{B}_k|, q)$ 
14   $H_k \leftarrow \text{SAMPLE}(\mathcal{B}_k, l)$                        //downsample new batch
15  if  $|H_k| > 0$  then                                     //insert new items
16     $S \leftarrow S \cup \{H_k\}$ 
17  output  $S$ 

```

Conceptually, at each time t_k , T-TBS first downsamples the current sample by independently flipping a coin for each item. The retention probability for an item depends on its age; specifically, an item $x \in \mathcal{B}_i$ is retained with probability $p_{i,k}$. T-TBS then downsamples the arriving batch \mathcal{B}_k via independent coin flips; an item in \mathcal{B}_k is inserted into the sample with probability q . As with B-TBS, the algorithm efficiently simulates multiple coin flips by directly generating the binomially distributed number of successes; thus the functions $\text{BINOMIAL}(j, r)$ and $\text{SAMPLE}(A, m)$ are defined as before.

Remark 1. The constraint that $b \geq n\gamma$ may lead to an inconveniently large required mean batch size. Intuitively, the problem is that an item's weight can become too small too quickly, even for polynomial decay. For instance, with $f_i = 1/(1+i)^2$, all items lose three fourths of their weight going from age 0 to age 1. For subexponential decay functions f , we can deal with this issue by using a shifted decay function $f_i^{(d)} = f_{d+i}/f_d$, where d is a positive integer. By choosing d sufficiently large, the corresponding value of $\gamma^{(d)} = f_0^{(d)}/\sum_{i=0}^{\infty} f_i^{(d)} = f_d/\sum_{i=d}^{\infty} f(i)$ can be made as small as desired. For instance, with $\Delta = 1$ and $f_i = 1/(1+i)^2$, we have $\gamma \approx 0.61$, whereas $\gamma^{(3)} \approx 0.22$. Of course, the original constraint requiring that $\Pr[x \in S_k]/\Pr[y \in S_k] = f_i/f_j$ for $x \in \mathcal{B}_i$ and $y \in \mathcal{B}_j$ is now modified to require that $\Pr[x \in S_k]/\Pr[y \in S_k] = f_{d+i}/f_{d+j}$, so that the relative inclusion probabilities have essentially the same “tail behavior” as f for large i and j , but the initial decay rate will be slower. This trick will not work for exponential decay, because here $f_i^{(d)} = e^{-\lambda d} f_i$, which implies that $\gamma^{(d)} \equiv \gamma$ for $d \geq 1$. In this case we must select λ small enough to accommodate the mean batch size. Thus non-exponential decay functions allow an additional degree of freedom when parameterizing the sampling algorithm. For superexponential decay, shifting will actually increase γ but, as discussed previously, such decay functions are of less practical interest. Over a broad range of experiments, quadratic decay with a shift of $d = 10$ yielded superior ML robustness results for both T-TBS and R-TBS, and we often use this variant in our experiments (Section 6).

3.2 Sample-Size Properties

We now analyze the sample size behavior of T-TBS, which directly impacts memory requirements, efficiency of memory usage, and ML model retraining time. We continue to assume that the batch sizes $\{|\mathcal{B}_k|\}_{k \geq 1}$ are i.i.d. with common mean $b \in [ny, \infty)$. Our first result (Theorem 3.1) describes the probabilistic behavior of the sample size $C_k = \sum_{H_i \in S_k} |H_i|$ for a fixed time t_k . Specifically, we give approximate expressions for the mean and variance of C_k when k is large. We also use Hoeffding's inequality to give exponential bounds valid for any k , showing that the probability of a very large deviation above or below the target value n at any given time t_k is very low. The proof of the theorem (and of most other results in the paper) is given in Appendix A. Denote by $\bar{b} \geq 1$ the maximum possible batch size, so that $\Pr[B \leq \bar{b}] = 1$. Recall that $F_k \triangleq \sum_{i=0}^k f_i$, and set $F_k^{(2)} = \sum_{i=0}^k f_i^2$.

THEOREM 3.1. *For any decay function f such that $F_\infty < \infty$,*

- (i) $\mathbb{E}[C_k] = nF_{k-1}/F_\infty \uparrow n$ as $k \rightarrow \infty$;
- (ii) $\text{Var}[C_k] \rightarrow bqF_\infty - bq^2F_\infty^{(2)}$;
- (iii) *if $\bar{b} < \infty$, then*
 - (a) $\Pr[C_k \geq (1 + \epsilon)n] \leq e^{-O(kn^2\epsilon^2)}$ for $\epsilon, k > 0$ and
 - (b) $\Pr[C_k \leq (1 - \epsilon)n] \leq e^{-O(kn^2)}$ for $\epsilon \in (0, 1)$ and sufficiently large k .

Thus, from (i), the expected sample size converges to the target size n as t becomes large and, from (ii), the variance also converges to a constant that depends on b and f . By (iii), the probability that the sample size deviates from n by more than $100\epsilon\%$ is exponentially small when k or n is large.

Remark 2. If f_i decays very slowly as $i \rightarrow \infty$, then the convergence of the expected sample size to n will also be very slow. For example, if $f_i = 1/(1 + i\Delta)^s$ for some $s > 1$, then, using (i) above and a standard bound, it is easy to show that $n - \mathbb{E}[C_k] = \Theta(1/k^{s-1})$. Thus choosing a value of, say, $s = 1.0001$ will result in a long sequence of undersized samples. Similarly, if the sample size becomes overly large at some point, recovery will be slow.

Theorem 3.1 does not tell the entire story. Although it follows from this theorem that, over many different sampling runs, the average sample size at a given (large) time t_k is close to n and the probability of being far away from n is small, the successive sample sizes during an individual sampling run need not be well behaved. This issue is addressed by Theorem 3.2 below. Assertion (i) shows that any sample size can be attained with positive probability, so one potential type of bad behavior might occur if, with positive probability, the sample size is unstable in that it drifts off to $+\infty$ over time. Assertion (ii) shows that such unstable behavior is ruled out if the maximum batch size is bounded and f_i decays rapidly enough so that equation (7) below holds. If f_i decays even faster, so that equation (8) below holds, then the stability assertion can be strengthened to guarantee that the times between successive attainments of a given sample size are not only all finite, but all have the same finite mean; moreover, the average sample size—averaged over times t_1, t_2, \dots, t_k —converges to n with probability 1 as k becomes large. On the negative side, it follows that, for a given sampling run, the sample size will repeatedly—though infrequently, since the expected sample size at any time point is finite—become arbitrarily large, even if the average behavior is good. This result shows that the sample-size control provided by T-TBS is incomplete, and thus motivates the more complex R-TBS algorithm.

In the following, write “i.o.” to denote that an event occurs “infinitely often”, i.e., for infinitely many values of k , and write “w.p.1” for “with probability 1”.

THEOREM 3.2. *Let the T-TBS decay function f satisfy $F_\infty < \infty$ and let \bar{b} be the maximum possible batch size. Then*

- (i) for all $m \geq 0$, there exists $k \geq 0$ such that $\Pr[C_k \geq m] > 0$;
- (ii) if $\bar{b} < \infty$ and

$$\sum_{i=0}^{\infty} i f_i < \infty, \quad (7)$$

then $\Pr[C_k = m \text{ i.o.}] = 1$ for all $m \geq 0$;

- (iii) if $\bar{b} < \infty$ and, for $k \geq 1$,

$$\sup_{i \geq 0} (f_{i+k} / f_i) \leq g_k \quad (8)$$

for some sequence $\{g_k\}_{k \geq 0}$ with $\lim_{k \rightarrow \infty} g_k = 0$, then (a) the expected times between successive visits to state m are uniformly bounded for any $m \geq 0$, and (b) $\lim_{k \rightarrow \infty} (1/k) \sum_{i=0}^k C_i = n$ w.p.1.

The proof of Theorem 3.2 rests on a reduced representation \mathcal{S}_k of the state S_k of the sample at a time t_k , comprising a collection of pairs of the form (n, i) , where n is the number of sample items of age Δi . In Appendix A we argue that the process $\{\mathcal{S}_k\}_{k \geq 0}$ is a time-homogeneous Markov chain, and hence we can use tools from the theory of Markov chains to establish the “recurrence” properties that correspond to our stability results. The state space of this Markov chain is quite complex, as are the transition probabilities between states, so the application of these tools is decidedly nontrivial.

Remark 3. Note that the assumptions on f indeed become increasingly strong when going from Assertions (i) to (iii). The condition in (7) trivially implies that $F_\infty < \infty$. Also, (8) implies (7). To see this, fix k large enough so that $g_k < 1$, and observe that, since $f_{i+k} \leq g_k f_i$ for all i ,

$$\sum_{i=0}^{\infty} i f_i \leq \sum_{m=0}^{k-1} f_m \left(\sum_{i=0}^{\infty} (ki + m) g_k^i \right) < \infty.$$

This increase in strength is strict: the decay function $f_i = 1/(i\Delta + 1)^2$ satisfies $F_\infty < \infty$ but not (7), and the decay function $f_i = 1/(i\Delta + 1)^3$ satisfies (7) but not (8). The condition in (7) holds, e.g., for exponential decay and for polynomial decay $f_i = 1/(1 + i\Delta)^s$ with $s > 2$. The condition in (8) holds, e.g., for functions that decay exponentially or faster.

Even in the most stable case, however, we do not have complete control over the sample size. Indeed, any sample size m , no matter how large, is exceeded infinitely often w.p.1 and the expected time between such incidents is uniformly bounded. Although the expected times are often very large, so that the incidents are infrequent, and the faster the decay, the faster the recovery from an incident, T-TBS is ultimately fragile with respect to sample size. This fragility is amplified when batch sizes fluctuate in a non-predictable way, as often happens in practice, and T-TBS can break down; see the experiments in Section 6.2.

Despite the fluctuations in sample size, T-TBS is of interest because, when the mean batch size is known and constant over time, and when some sample overflows are tolerable, T-TBS is relatively simple to implement and parallelize, and is very fast (see Section 6). For example, if the data comes from periodic polling of a set of robust sensors, the data arrival rate will be known a priori and will be relatively constant, except for the occasional sensor failure, and hence T-TBS might be appropriate. On the other hand, if data is coming from, e.g., a social network, then batch sizes may be hard to predict.

4 RESERVOIR-BASED TBS

Targeted time-biased sampling (T-TBS) controls the decay rate but only partially controls the sample size, whereas batched reservoir sampling (B-RS) bounds the sample size but does not allow time biasing. Our new reservoir-based time-biased sampling algorithm (R-TBS) combines the best

features of both, controlling the decay rate while ensuring that the sample never overflows. For exponential decay, R-TBS has optimal sample size and stability properties, and in the general case the user can trade off storage for both sample-size stability and accuracy. Importantly, unlike T-TBS, the R-TBS algorithm can handle any sequence of batch sizes.

4.1 Item Weights and Latent Samples

R-TBS combines the use of a reservoir with the notion of “latent samples” to enforce (1) and bound the sample size. Latent samples, in turn, rest upon the notion of “item weights”.

Item weights: In R-TBS, the *weight* of an item of age α is given by $f(\alpha)$, where f is the decay function; note that a newly arrived item has a weight of $f(0) = 1$. As discussed later, R-TBS ensures that the probability that an item appears in the sample is proportional (or approximately proportional) to its weight. All items arriving at the same time have the same weight, so that the *total weight* of all items seen up through time t_k is $W_k = \sum_{i=1}^k |\mathcal{B}_i| f(\alpha_{i,k})$. For traditional (sequential or batch-oriented) reservoir sampling, an item does not decay, and so has a weight equal to 1 at all times; thus the notions of items and item weights coincide. Moreover, in the traditional setting the weight of a sample coincides with the number of items in the sample. In our generalized setting, the “size” (weight) of a sample and the number of items in the sample differ, with samples having fractional sizes. We handle this complication via the notion of a *latent fractional sample*.

ALGORITHM 4: Generating a sample from a latent sample

```

1  $L = (A, \pi, C)$ : latent sample
2  $U \leftarrow \text{UNIFORM}()$ 
3 if  $U \leq \text{frac}(C)$  then  $S \leftarrow A \cup \pi$  else  $S \leftarrow A$ 
4 return  $S$ 

```

Latent samples: A latent fractional sample formalizes the idea of a sample of fractional size. Formally, given a set U of items, a *latent sample* of U with *sample weight* C is a triple $L = (A, \pi, C)$, where $A \subseteq U$ is a set of $\lfloor C \rfloor$ *full* items and $\pi \subseteq U$ is a (possibly empty) set containing at most one *partial* item; π is nonempty if and only if $C > \lfloor C \rfloor$.




Fig. 1. Latent sample L (sample weight $C = 3.6$) and possible realized samples

R-TBS maintains a latent sample L over time and produces an actual sample S from L on demand by sampling as described in Algorithm 4; see Figure 1 for an example. In the pseudocode, $\text{frac}(x) = x - \lfloor x \rfloor$ and the function $\text{UNIFORM}()$ generates a random number uniformly distributed on $[0, 1]$. Because each full item is included with probability 1 and the partial item is included with probability $\text{frac}(C)$, we have

$$\mathbb{E}[|S|] = \lceil C \rceil \text{frac}(C) + \lfloor C \rfloor (1 - \text{frac}(C)) = (\lceil C \rceil - \lfloor C \rfloor) \text{frac}(C) + \lfloor C \rfloor = \text{frac}(C) + \lfloor C \rfloor = C, \quad (9)$$

so that the size of S equals C in expectation. By allowing at most one partial item, we minimize the latent sample’s footprint: $|A \cup \pi| \leq \lfloor C \rfloor + 1$. Importantly, if the weight C of a latent sample L

ALGORITHM 5: Downsampling

```

1  $L = (A, \pi, C)$ : input latent sample
2  $\theta$ : scaling factor with  $\theta \in (0, 1)$ 
3  $U \leftarrow \text{UNIFORM}()$ ;  $C' = \theta C$ 
4 if  $\lfloor C' \rfloor = 0$  then                                     //no full items retained
5     if  $U > \text{frac}(C)/C$  then
6          $(A', \pi') \leftarrow \text{SWAP1}(A, \pi)$ 
7          $A' \leftarrow \emptyset$ 
8 else if  $0 < \lfloor C' \rfloor = \lfloor C \rfloor$  then                     //no items deleted
9     if  $U > (1 - \theta \text{frac}(C))/(1 - \text{frac}(C'))$  then
10         $(A', \pi') \leftarrow \text{SWAP1}(A, \pi)$ 
11 else                                                     //items deleted:  $0 < \lfloor C' \rfloor < \lfloor C \rfloor$ 
12     if  $U \leq \theta \text{frac}(C)$  then
13          $A' \leftarrow \text{SAMPLE}(A, \lfloor C' \rfloor)$ 
14          $(A', \pi') \leftarrow \text{SWAP1}(A', \pi)$ 
15     else
16          $A' \leftarrow \text{SAMPLE}(A, \lfloor C' \rfloor + 1)$ 
17          $(A', \pi') \leftarrow \text{MOVE1}(A', \pi)$ 
18 if  $C' = \lfloor C' \rfloor$  then                                     //no fractional item
19      $\pi' \leftarrow \emptyset$ 
20 return  $L' = (A', \pi', C')$ 

```

is an integer, then L contains no partial item, and the sample S generated from L via Algorithm 4 is unique and contains exactly C items; thus, the sample weight and the number of sample items coincide in this case. We now describe two key operations on latent samples that are used by R-TBS.

Downsampling: Besides extracting an actual sample from a latent sample, another key operation on latent samples is *downsampling*. For $\theta \in (0, 1)$, the goal of downsampling $L = (A, \pi, C)$ is to obtain a new latent sample $L' = (A', \pi', \theta C)$ such that, if we generate S and S' from C and C' via Algorithm 4, we have

$$\Pr[x \in S'] = \theta \Pr[x \in S] \quad (10)$$

for all $x \in S$. Thus the appearance probability for each item in S , as well as the expected size of the sample $E[S]$, is scaled down by a factor of θ . Theorem 4.1 (later in this section) asserts that Algorithm 5 satisfies this property.

In the pseudocode for Algorithm 5, the subroutine $\text{SWAP1}(A, \pi)$ moves a randomly selected item from A to π and moves the current item in π (if any) to A . Similarly, $\text{MOVE1}(A, \pi)$ moves a randomly selected item from A to π , replacing the current item in π (if any). More precisely, $\text{SWAP1}(A, \pi)$ executes the operations $I \leftarrow \text{SAMPLE}(A, 1)$, $A \leftarrow (A \setminus I) \cup \pi$, and $\pi \leftarrow I$, and $\text{MOVE1}(A, \pi)$ executes the operations $I \leftarrow \text{SAMPLE}(A, 1)$, $A \leftarrow A \setminus I$, and $\pi \leftarrow I$.

To gain some intuition for why the algorithm works, consider a simple special case, where the goal is to form a latent sample $L' = (A', \pi', \theta C)$ from a latent sample $L = (A, \pi, C)$ of integral size C ; that is, L comprises exactly C full items. Assume that $C' = \theta C$ is non-integral, so that L' contains a partial item, and that $\lfloor C \rfloor > \lfloor C' \rfloor$; e.g., $C = 3$ and $C' = 1.5$, so that $\theta = 0.5$. In this case, we simply select an item at random (from A) to be the partial item in L' and then select $\lfloor C' \rfloor$ of the remaining $C - 1$ items at random to be the full items in L' ; see Figure 2(a). Denote by S and S' the samples obtained from L and L' via Algorithm 4. By symmetry, each item $i \in L$ is equally likely to be included in S' , so that the inclusion probabilities for the items in L are all scaled down by the same fraction, as required for (10). In Figure 2(a), for example, item a appears in S with probability 1 since

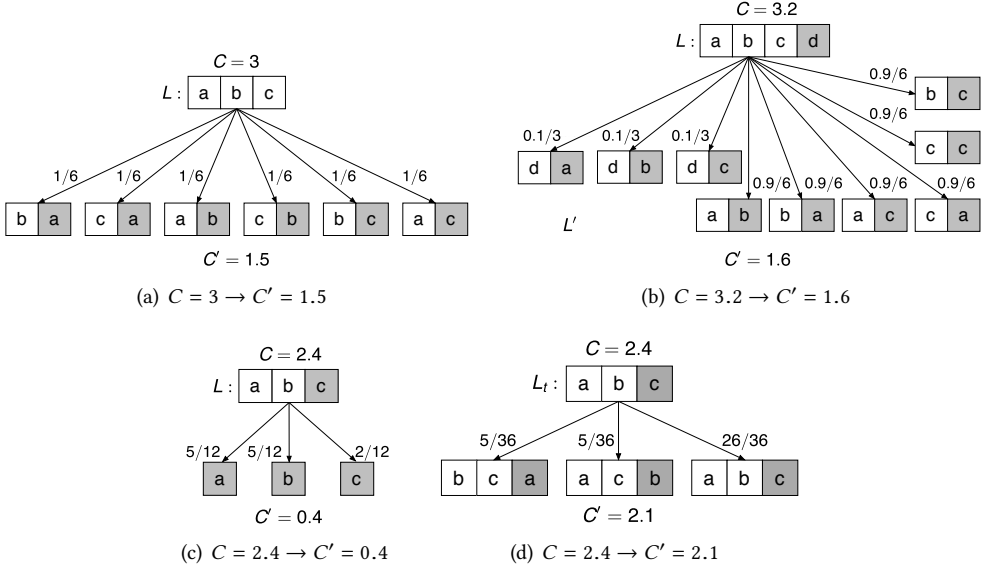


Fig. 2. Downsampling examples

it is a full item. In S' , where the weights have been reduced by 50%, item a (either as a full or partial item, depending on the random outcome) appears with probability $2 \cdot (1/6) + 2 \cdot (1/6) \cdot 0.5 = 0.5$, as expected. This scenario corresponds to lines 16 and 17 in the algorithm, where we carry out the above selections by randomly sampling $\lfloor C' \rfloor + 1$ items from A to form A' and then choosing a random item in A' as the partial item by moving it to π .

In the case where L contains a partial item x^* that appears in S with probability $\text{frac}(C)$, it follows from (10) that x^* should appear in S' with probability $p = \theta P[x^* \in S] = \theta \text{frac}(C)$. Thus, with probability p , lines 13–14 retain x^* and convert it to a full item so that it appears in S' . Otherwise, in lines 16–17, x^* is removed from the sample when it is overwritten by a random item from A' ; see Figure 2(b). Again, a new partial item is chosen from A in a random manner to uniformly scale down the inclusion probabilities. For instance, in Figure 2(b), item d appears in S with probability 0.2 (because it is a partial item) and in S' , appears with probability $3 \cdot (0.1/3) = 0.1$. Similarly, item a appears in S with probability 1 and in S' with probability $(1.8)/6 + 0.6 \cdot (1.8/6) + 0.6 \cdot (0.1/3) = 0.5$.

The if-statement in line 4 corresponds to the corner case in which L' does not contain a full item. The partial item $x^* \in L$ either becomes full or is swapped into A' and then immediately ejected; see Figure 2(c).

The if-statement in line 8 corresponds to the case in which no items are deleted from the latent sample, e.g., when $C = 4.7$ and $C' = 4.2$. In this case, x^* either becomes full by being swapped into A' or remains as the partial item for L' . Denoting by ρ the probability of *not* swapping, we have $P[x^* \in S'] = \rho \cdot \text{frac}(C') + (1 - \rho) \cdot 1$. On the other hand, (10) implies that $P[x^* \in S'] = \theta \text{frac}(C)$. Equating these expressions shows that ρ must equal the expression on the right side of the inequality on line 9; see Figure 2(d).

THEOREM 4.1. *For $\theta \in (0, 1)$, let $L' = (A', \pi', \theta C)$ be the latent sample produced from a latent sample $L = (A, \pi, C)$ via Algorithm 5, and let S' and S be samples produced from L' and L via Algorithm 4. Then $\Pr[x \in S'] = \theta \Pr[x \in S]$ for all $x \in A \cup \pi$.*

The union operator: We also need to take the union of disjoint latent samples while preserving the inclusion probabilities for each. Two latent samples $L_1 = (A_1, \pi_1, C_1)$ and $L_2 = (A_2, \pi_2, C_2)$ are *disjoint* if $(A_1 \cup \pi_1) \cap (A_2 \cup \pi_2) = \emptyset$. The pseudocode for the union operation is given as Algorithm 6. The idea is to add all full items to the combined latent sample. If there are partials items in L_1 and L_2 , then we transform them to either a single partial item, a full item, or a full plus partial item, depending on the values of $\text{frac}(C_1)$ and $\text{frac}(C_2)$. Such transformations are done in a manner that preserves the appearance probabilities. Of course, we can obtain the union of an arbitrary number of latent samples by iterating Algorithm 6; for latent samples L_1, \dots, L_k , we denote by $\bigcup_{j=1}^k L_j$ the latent sample produced by this procedure.

ALGORITHM 6: Union

```

1  $L_1 = (A_1, \pi_1, C_1)$ : fractional sample of size  $C_1$ 
2  $L_2 = (A_2, \pi_2, C_2)$ : fractional sample of size  $C_2$ 
3  $C \leftarrow C_1 + C_2$ 
4  $U \leftarrow \text{UNIFORM}()$ 
5 if  $\text{frac}(C_1) + \text{frac}(C_2) < 1$  then
6    $A \leftarrow A_1 \cup A_2$ 
7   if  $U \leq \text{frac}(C_1) / (\text{frac}(C_1) + \text{frac}(C_2))$  then  $\pi \leftarrow \pi_1$  else  $\pi \leftarrow \pi_2$ 
8 else if  $\text{frac}(C_1) + \text{frac}(C_2) = 1$  then
9    $\pi \leftarrow \emptyset$ 
10  if  $U \leq \text{frac}(C_1)$  then  $A \leftarrow A_1 \cup A_2 \cup \pi_1$  else  $A \leftarrow A_1 \cup A_2 \cup \pi_2$ 
11 else  $\text{frac}(C_1) + \text{frac}(C_2) > 1$ 
12   if  $U \leq (1 - \text{frac}(C_1)) / [(1 - \text{frac}(C_1)) + (1 - \text{frac}(C_2))]$  then
13      $\pi = \pi_1$ 
14      $A \leftarrow A_1 \cup A_2 \cup \pi_2$ 
15   else
16      $\pi = \pi_2$ 
17      $A \leftarrow A_1 \cup A_2 \cup \pi_1$ 
18 return  $L = (A, \pi, C)$ 

```

THEOREM 4.2. Let $L_1 = (A_1, \pi_1, C_1)$ and $L_2 = (A_2, \pi_2, C_2)$, be disjoint latent samples, and let $L = (A, \pi, C)$ be the latent sample produced from L_1 and L_2 by Algorithm 6. Let S_1, S_2 , and S be random samples generated from L_1 , and L_2 , and L via Algorithm 4. Then

- (i) $C = C_1 + C_2 = E[S]$;
- (ii) $\forall x \in L_1, \Pr[x \in S] = \Pr[x \in S_1]$; and
- (iii) $\forall x \in L_2, \Pr[x \in S] = \Pr[x \in S_2]$.

4.2 The R-TBS Algorithm with Exponential Decay

Our general goal is to provide a sampling algorithm that bounds the sample size at n while enforcing (1). For the special case of exponential decay, this task is greatly facilitated by the fact that, at each time step, all items in the sample decay by the same multiplicative factor. We exploit this fact to provide a relatively simple version of R-TBS for the case of exponential decay. In Section 4.3, we show how to generalize our approach to the case of arbitrary decay functions.

The algorithm: R-TBS for exponential decay is given as Algorithm 7. The algorithm generates a sequence of latent samples $\{L_k\}_{k \geq 1}$ and from these generates a sequence of actual samples $\{S_k\}_{k \geq 1}$ that are returned to the user. In the algorithm, the functions GETSAMPLE, DOWNSAMPLE, and UNION execute the operations described in Algorithms 4, 5, and 6.

ALGORITHM 7: Reservoir-based TBS (R-TBS) for exponential decay

```

1  $\lambda$ : decay rate ( $\geq 0$ )
2  $n$ : maximum sample size
3 Initialize:  $W \leftarrow 0$ ;  $A \leftarrow \emptyset$ ;  $\pi \leftarrow \emptyset$ ;  $C \leftarrow 0$ ;  $t_0 \leftarrow 0$ ;  $\rho \leftarrow 1$ 
4 for  $k \leftarrow 1, 2, \dots$  do
5    $\theta \leftarrow e^{-\lambda(t_k - t_{k-1})}$  //decay factor
6    $W \leftarrow \theta W + |\mathcal{B}_k|$  //update total weight
7    $\rho' \leftarrow \rho$ 
8    $\rho \leftarrow \min(1, n/W)$ 
9   if  $C > 0$  then  $(A, \pi, C) \leftarrow \text{DOWNSAMPLE}((A, \pi, C), (\rho/\rho')\theta)$  //decay old items
10   $L_0 \leftarrow \text{DOWNSAMPLE}((\mathcal{B}_k, \emptyset, |\mathcal{B}_k|), \rho)$  //take in new items
11   $L \leftarrow \text{UNION}(L_0, (A, \pi, C))$  //combine old and new items
12   $S \leftarrow \text{GETSAMPLE}(L)$ 
13  output  $S$ 

```

The goal of the algorithm is to ensure that

$$\Pr[x \in S_k] = \rho_k f(\alpha_{i,k}) \quad (11)$$

for all $k \geq 1$, $i \leq k$, and $x \in \mathcal{B}_i$, where $f(\alpha) = e^{-\lambda\alpha}$ and $\{\rho_k\}_{k \geq 1}$ are the successive values of the variable ρ during a run of the algorithm. Clearly, (11) immediately implies (1). We choose ρ_k to make the sample size as large as possible without exceeding n . In more detail, we show in Theorem 4.3 below that $C_k = \rho_k W_k$ for all k . We therefore set $\rho_k = \min(1, n/W_k)$ —see line 8—so that $C_k = \min(W_k, n)$. Thus if $W_k < n$, then the sample weight is at its maximum possible value W_k , leading to the maximum possible sample size of either $\lfloor W_k \rfloor$ or $\lceil W_k \rceil$. If $W_k \geq n$, then the sample weight, and hence the sample size, is capped at n . The algorithm functions analogously to classic reservoir sampling: if the (weighted) items seen so far can fit into the reservoir of size n , then they are simply accepted, if the total item weight exceeds n , then, when a new batch arrives, a random subset of old items is removed from the sample via downsampling (line 9) and a random subset of the arriving items, also filtered via downsampling (line 10), take their place (line 11). Note that if $|\mathcal{B}_j| \equiv 1$ for all j , so that we process items one at a time, and if there is no decay, so that $f(\alpha) \equiv 1$, then $W_k = k$ and the inclusion probability in (11) reduces to $\Pr[x \in S_k] = \rho_k \cdot 1 = q_k$, where $q_k = \min(1, n/k)$, exactly as in traditional reservoir sampling.

Algorithm properties: Theorem 4.3(i) below asserts that R-TBS satisfies (11) and hence (1), thereby maintaining the correct inclusion probabilities. Indeed, suppose that the inclusion probability for $x \in \mathcal{B}_i$ (with $i < k$) at time t_{k-1} is $\rho_{k-1} f(\alpha_{i,k-1})$. Write $\theta_k = e^{-\lambda(t_k - t_{k-1})}$ and observe that $\theta_k = f(\alpha_{i,k})/f(\alpha_{i,k-1})$ for any $i \in [0..k-1]$. It follows that

$$\rho_{k-1} f(\alpha_{i,k-1}) \cdot \left(\frac{\rho_k}{\rho_{k-1}} \right) \theta_k = \rho_{k-1} f(\alpha_{i,k-1}) \cdot \left(\frac{\rho_k f(\alpha_{i,k})}{\rho_{k-1} f(\alpha_{i,k-1})} \right) = \rho_k f(\alpha_{i,k}),$$

preserving the desired inclusion probability (11); the downsampling operation in line 9 executes this adjustment. (See Appendix A for the detailed inductive proof.) Similarly, an incoming item $x \in \mathcal{B}_k$ is accepted into the sample with probability $\rho_k \cdot 1 = \rho_k f(\alpha_{k,k})$ —see line 10—so that (11) holds in this case as well. Note that we can combine the above results over all of the batches by virtue of Theorem 4.2. Theorem 4.3(ii) implies that the sample size and stability are maximized, as formalized in Theorems 4.4 and 4.5 below. Finally, the assertion in Theorem 4.3(iii) ensures that the inclusion probabilities for a given item are nonincreasing over time. This is crucial, since otherwise we might have to recover an item that was previously deleted from the sample, which is impossible. This monotonicity property trivially holds for traditional one-item-at-a-time reservoir

sampling, where $\rho_k f(\alpha_{i,k}) = \min(1, n/k)$ as discussed previously. For general decay functions, the monotonicity property hinges on the interplay of item decay and new-item arrival, and in fact does not generally hold if we try to define ρ_k simply as $\min(1, n/W_k)$. To deal with this issue, we need a more complex scheme for defining ρ_k ; see Section 4.3.3.

THEOREM 4.3. *Let $\{L_k = (A_k, \pi_k, C_k)\}_{k \geq 1}$ and $\{S_k\}_{k \geq 1}$ be a sequence of latent samples and samples, respectively, produced by Algorithm 7 and define $\rho_k = \min(1, n/W_k)$. Then*

- (i) $\Pr[x \in S_k] = \rho_k f(\alpha_{i,k})$ for all $k \geq 1$, $i \leq k$, and $x \in \mathcal{B}_i$;
- (ii) $C_k = \rho_k W_k$ for all k ; and
- (iii) $\rho_k f(\alpha_{i,k}) \leq \rho_{k-1} f(\alpha_{i,k-1})$ for all $k > 1$ and $i < k$.

We call a sample S_k *unsaturated* if $C_k < n$ and *saturated* if $C_k = |S_k| = n$; note that we also have $W_k < n$ if and only if the sample is unsaturated. Theorems 4.4 and 4.5 below assert that, among all sampling algorithms with exponential time biasing, R-TBS both maximizes the expected sample size in unsaturated scenarios and minimizes sample-size variability. Thus R-TBS tends to yield more accurate ML results (via more training data) and greater stability in both result quality and retraining costs.

THEOREM 4.4. *Let H be any sampling algorithm for exponential decay that satisfies (1) and denote by S_k and S_k^H the samples produced at time t_k by R-TBS and H . If the total weight at some time $t_k \geq 1$ satisfies $W_k < n$, then $\mathbb{E}[|S_k^H|] \leq \mathbb{E}[|S_k|]$.*

PROOF. Since H satisfies (1), it follows that, for each time $t_j \leq t_k$ and $x \in \mathcal{B}_j$, the inclusion probability $\Pr[x \in S_k^H]$ must be of the form $r_k e^{-\lambda(t_k - t_j)}$ for some function r_k independent of j . Taking $j = k$, we see that $r_k \leq 1$. For R-TBS in an unsaturated state, (11) implies that $r_k = \rho_k = C_k/W_k = 1$, so that $\Pr[x \in S_k^H] \leq \Pr[x \in S_k]$, and the desired result follows directly. \square

THEOREM 4.5. *Let H be any sampling algorithm for exponential decay that satisfies (1) and has maximal expected sample size C_k , and denote by S_k and S_k^H the samples produced at time t_k by R-TBS and H . Then $\text{Var}[|S_k^H|] \geq \text{Var}[|S_k|]$ for any $k \geq 1$.*

PROOF. Considering all possible distributions over the sample size having a mean value equal to C_k , it is straightforward to show that variance is minimized by concentrating all of the probability mass onto $\lfloor C_k \rfloor$ and $\lceil C_k \rceil$. There is precisely one such distribution, which results from application of Algorithm 4, and this is precisely the sample-size distribution attained by R-TBS. \square

4.3 The General R-TBS Algorithm

For a general decay function, the decay factor for items in different batches is no longer the same, as in the exponential case, which adds substantial complexity to the R-TBS algorithm. In particular, we need to track the timestamps of individual items and so, analogously to T-TBS, we represent the state \mathbb{L} of the sample as a set of triples of the form $L_i = (A_i, \pi_i, C_i)$, where L_i is a latent sample of items from \mathcal{B}_i . For $k \geq 1$ we denote by \mathbb{L}_k the state of the sample at time t_k and by $\mathbb{C}_k = \sum_{(A_i, \pi_i, C_i) \in \mathbb{L}_k} |C_i|$ the total sample weight at time t_k . The pseudocode for the general R-TBS procedure is given as Algorithm 8; the ideas behind the algorithm are developed below.

4.3.1 A naive algorithm. The core idea of Algorithm 8 is to try and mimic the exponential case. The naive version of doing this is as follows. When batch \mathcal{B}_k arrives, we first update the total weight W_{k-1} by updating the weight $W_{k-1,i}$ of each batch \mathcal{B}_i (line 9) for $i < k$, and then adding the weight of batch \mathcal{B}_k (line 11), which is just $|\mathcal{B}_k|$ since each incoming item has weight $f(0) = 1$. (By the “weight of batch \mathcal{B}_j ,” we mean the total weight of all items in \mathcal{B}_j .) Note that arguments

ALGORITHM 8: Reservoir-based TBS (R-TBS)

```

1  $f$ : decay function    $n$ : maximum sample size    $n'$ : maximum sample weight    $\lambda$ : decay rate for consolidated sample
2  $\delta_1 \in (0, 1)$  and  $\delta_2 > 0$ : approximation parameters for consolidated latent sample
3  $m$ : dynamic parameter such that items arriving at time  $t_i < t_m$  are in consolidated latent sample
4 Initialize:  $(\tilde{A}, \tilde{\pi}, \tilde{C}) \leftarrow (\emptyset, \emptyset, 0)$ ;  $\mathbb{L} \leftarrow \{(\tilde{A}, \tilde{\pi}, \tilde{C})\}$ ;  $W \leftarrow 0$ ;  $F_\infty = 1/\text{GAMMA}(f)$ ;  $\rho \leftarrow 1$ ;  $B^* \leftarrow 0$ ;  $m \leftarrow 1$ 
5 for  $k \leftarrow 1, 2, \dots$  do
6    $B^* \leftarrow \max(B^*, |\mathcal{B}_k|)$ 
7   //update total weight
8   if  $k > 1$  then
9     for  $i \leftarrow k-1, k-2, \dots, m$  do                                     //decay weight of recent items
10       $W \leftarrow W - \left(1 - \frac{f(\alpha_{i,k})}{f(\alpha_{i,k-1})}\right)(C_i/\rho)$ 
11       $W \leftarrow W - (1 - e^{-\lambda\Delta})(\tilde{C}/\rho)$                                      //decay weight of consolidated items
12     $W \leftarrow W + |\mathcal{B}_k|$                                                  //add weight of new items
13    //update  $\rho$ 
14    if  $k > 1$  then
15       $\rho' \leftarrow \rho$ 
16       $\rho^* \leftarrow \min_{m \leq i \leq k} \rho' f(\alpha_{i,k-1})/f(\alpha_{i,k})$ 
17      if  $m > 1$  then  $\rho^* \leftarrow \min(\rho^*, \rho' e^{\lambda\Delta})$ 
18       $\rho \leftarrow \min(1, n'/W, \rho^*)$ 
19    //update samples
20    if  $k > 1$  then
21      for  $i \leftarrow k-1, k-2, \dots, m$  do                                     //update recent samples
22         $(A_i, \pi_i, C_i) \leftarrow \text{DOWNSAMPLE}((A_i, \pi_i, C_i), \frac{\rho f(\alpha_{i,k})}{\rho' f(\alpha_{i,k-1})})$ 
23         $(\tilde{A}, \tilde{\pi}, \tilde{C}) \leftarrow \text{DOWNSAMPLE}((\tilde{A}, \tilde{\pi}, \tilde{C}), (\rho/\rho')e^{-\lambda\Delta})$        //update consolidated latent sample
24       $(A_k, \pi_k, C_k) \leftarrow \text{DOWNSAMPLE}((\mathcal{B}_k, \emptyset, |\mathcal{B}_k|), \rho)$ 
25       $\mathbb{L} \leftarrow \mathbb{L} \cup \{(A_k, \pi_k, C_k)\}$                                      //add sample of new items
26      //try to consolidate samples
27      while  $(f(\alpha_{m,k}) < \delta_1) \wedge (F_\infty - \sum_{i=m}^k f(\alpha_{i,k}) < \delta_2/B^*)$  do
28         $(\tilde{A}, \tilde{\pi}, \tilde{C}) \leftarrow \text{UNION}((\tilde{A}, \tilde{\pi}, \tilde{C}), (A_m, \pi_m, C_m))$ 
29         $\mathbb{L} \leftarrow \mathbb{L} \setminus \{(A_m, \pi_m, C_m)\}$ 
30         $m \leftarrow m + 1$ 
31      //output sample to user
32       $(A, \pi, C) \leftarrow \text{UNION}((A_k, \pi_k, C_k), (A_{k-1}, \pi_{k-1}, C_{k-1}), \dots, (A_m, \pi_m, C_m), (\tilde{A}, \tilde{\pi}, \tilde{C}))$ 
33      if  $C > n$  then  $(A, \pi, C) \leftarrow \text{DOWNSAMPLE}((A, \pi, C), n)$ 
34       $S \leftarrow \text{GETSAMPLE}((A, \pi, C))$ 
35    output  $S$ 

```

essentially identical to those in Theorem 4.3 show that $C_i = \rho_j W_{j,i}$ for all $j \geq 1$ and $i \leq j$, where ρ_j is defined appropriately (see Section 4.3.3 below). Thus the term C_i/ρ in line 9 is precisely $W_{k-1,i}$, which loses a fraction $\rho_k f(\alpha_{i,k})/(\rho_{k-1} f(\alpha_{i,k-1}))$ of its weight. We next downsample each latent sample L_i by a factor of $\rho_k f(\alpha_{i,k})/(\rho_{k-1} f(\alpha_{i,k-1}))$ for $i < k$ —see line 19—and then downsample $(\mathcal{B}_k, \emptyset, |\mathcal{B}_k|)$ by a factor of ρ_k (line 21) to create L_k . To output a sample \mathbb{S}_k to the user, we first union L_1, \dots, L_k using Algorithm 6 (line 27) and then create \mathbb{S}_k via Algorithm 4 (line 29).

The algorithm as described indeed satisfies (1). Specifically, denoting by S_i a sample created from L_i via Algorithm 4, arguments as before show that $\Pr[x \in S_i] = \rho_k f(\alpha_{i,k})$ for $x \in \mathcal{B}_i$ with $i \leq k$. Then Theorem 4.2 implies that $\Pr[x \in \mathbb{S}_k] = \rho_k f(\alpha_{i,k})$ for all $i \leq k$ and $x \in \mathcal{B}_i$, so that (1)

holds. Moreover, $\mathbb{C}_k = \rho_k W_k$ as before. This naive algorithm, however, has two issues that take some effort to address.

4.3.2 Sample footprint. Perhaps the most important problem with the naive algorithm is that the sample footprint grows without bound. To see this, observe that the latent sample for items in a given batch \mathcal{B}_i never empties out completely. At time goes on, the latent sample will eventually contain one partial item, whose appearance probability is always positive (though decreasing to 0). Thus the sample footprint at time t_k is $\Omega(k)$.

Our solution to this problem is to approximate the exact time biasing scheme in the naive algorithm by maintaining, at each time t_k , distinct latent samples only for items that have arrived at time $t_{m(k)}$ or later, where $m(k)$ is a carefully chosen index that increases with k . We denote by $\alpha_k^* = \alpha_{m(k),k}$ the age at time t_k of items that arrived at $t_{m(k)}$. Sample items that arrived earlier than time $t_{m(k)}$, i.e., whose age is greater than α_k^* , are maintained in a single *consolidated* latent sample, which decays at an exponential rate λ . Thus \mathbb{L}_k comprises $k - m(k) + 2$ latent samples in total. The values of $m(k)$ and λ are determined by parameters δ_1 and δ_2 that control the accuracy of the approximate time biasing scheme, as described below.

Recall our running assumption that $t_i = i\Delta$ for $i \geq 0$ and some $\Delta > 0$. We also assume that $F_\infty < \infty$ —where $f_i = f(i\Delta)$ and $F_k = \sum_{j=0}^k f_j$ as before—and that λ is chosen to ensure that

$$e^{-\lambda\Delta} \leq f(\alpha + \Delta)/f(\alpha) \quad (12)$$

for all α large enough so that $f(\alpha) < \delta_1$. This is always possible for subexponential decay functions. For example, if $f(\alpha) = 1/(1 + \alpha)^s$, then (12) holds for any $\lambda \geq s \ln((1 + \hat{\alpha} + \Delta)/(1 + \hat{\alpha}))/\Delta$, where $\hat{\alpha}$ is the smallest value of α such that $f(\alpha) < \delta_1$. We also assume that $\bar{b} \triangleq \sup_i |\mathcal{B}_i| < \infty$.

The use of a consolidated sample is equivalent to using, at each time t_k , a modified decay function \tilde{f}_k given by

$$\tilde{f}_k(\alpha) = \begin{cases} f(\alpha) & \text{if } \alpha \leq \alpha_k^*; \\ f(\alpha_k^*)e^{-\lambda(\alpha - \alpha_k^*)} & \text{if } \alpha > \alpha_k^*. \end{cases}$$

Thus the appearance probability for recent items is governed exactly by the desired function f , whereas the appearance probability for older items is perturbed. Set $N = \min\{n \geq 1 : \sum_{i=n}^\infty f_i \leq \delta_2/\bar{b}\}$ and observe that $N < \infty$ by finiteness of F_∞ . Theorem 4.6 below shows that, at each time t_k , (i) the parameter δ_1 in Algorithm 8 is a bound on the absolute difference between f and \tilde{f}_k , i.e., the amount by which an older item's appearance probability is perturbed, (ii) the parameter δ_2 is a bound on the expected number of older items whose appearance probability is perturbed, and (iii) at most $N + 2$ latent samples need to be stored at any time point. This bound on the latent samples, coupled with the bound on the sample weights enforced by the reservoir capacity, ensures that the sample footprint is bounded.

THEOREM 4.6. *Algorithm 8 ensures the following properties for each t_k :*

- (i) $|f(\alpha) - \tilde{f}_k(\alpha)| \leq \delta_1$ for all $\alpha \geq \alpha_k^*$;
- (ii) $\sum_{i=1}^{m(k)-1} |\mathcal{B}_i| f(\alpha_{i,k}) < \delta_2$; and
- (iii) $k - m(k) \leq N$.

Examination of the proof of this theorem shows that the key invariants that must be maintained are (i) $f(\alpha) < \delta_1$ for $\alpha > \alpha_k^*$ and (ii) $F_\infty - \sum_{i=m(k)}^k f(\alpha_{i,k}) < \delta_2/\max_{1 \leq i \leq k} |\mathcal{B}_i|$. Lines 23–26 maintain these invariants while trying to add as many latent samples as possible to the consolidated sample. Note that the $\text{GAMMA}(f)$ function is the same function that is used to compute $1/F_\infty$ in the T-TBS algorithm. The code in lines 10, 15, and 20 performs the same operations on the consolidated latent sample as the code in the adjacent lines performs on the other latent samples.

Note that N is typically a conservative upper bound, both because of the use of the conservative constant \bar{b} and the fact that we have implicitly used the upper bound of $\rho_k = 1$ in our analysis for $k \geq 1$. Suppose, for example, that $\Delta = 1$, $\bar{b} = 10,000$, $n = 100,000$, and $f(\alpha) = 1/(1 + \alpha)^2$. If $\delta_1 = 10^{-4}$ and $\delta_2 = 100$ items—so that only 0.1% of sample items have perturbed appearance probabilities—then the number of latent samples stored is bounded above by $N + 2 \approx 100$. This corresponds to storage of up to 100 partial items, and represents about a 0.1% storage overhead relative to exponential decay (which stores at most one partial item). When f decays slowly, however, the overhead can become substantial. Thus, as in T-TBS, very slowly decaying functions can be expensive and cumbersome from a practical standpoint.

4.3.3 Monotonicity of inclusion probabilities. The remaining issue is that, as hinted above, we cannot simply choose $\rho_k = \min(1, n/W_k)$, even though this would produce the largest possible sample sizes. The problem is that, with this choice, the resulting appearance probability $p_{i,k} = \rho_k f(\alpha_{i,k}) = \min(1, n/W_k) f(\alpha_{i,k})$ is no longer guaranteed to be nonincreasing in k . This would cause the algorithm to break. In particular, the downsizing operation for latent sample L_i would actually try to *upsample* the sample, which is impossible since any \mathcal{B}_i items not in the sample have been discarded. For example, suppose that $n = 1000$ and $W_{k-1} = 2000$, and that $|\mathcal{B}_i| = 100$ and $f(\alpha_{i,k-1}) = 0.1$ for some batch \mathcal{B}_i with $i < k$. Then $p_{i,k-1} = \rho_{k-1} f(\alpha_{i,k-1}) = (n/W_{k-1}) f(\alpha_{i,k-1}) = 0.5 \cdot 0.1 = 0.05$, so that the expected number of sample items from \mathcal{B}_i at time t_{k-1} is $|\mathcal{B}_i| \cdot 0.05 = 5$. Now suppose that $f(\alpha_{i,k})/f(\alpha_{i,k-1}) = 0.9$ so that $f(\alpha_{i,k}) = 0.09$, but that W_k drops to 1000 through rapid decay in other batches. Then $\rho_k = (n/W_k) = 1$ and hence $p_{i,k} = 1 \cdot 0.09 = 0.09 > p_{i,k-1}$. The largest inclusion probability we can support is $\rho^* f(\alpha_{i,k}) = 0.05$, where $\rho^* = \rho_{k-1} f(\alpha_{i,k-1})/f(\alpha_{i,k}) = 5/9$. Our problem arises because the large 50% decrease in the total weight causes a factor of 2 increase in ρ , which overwhelms the factor of 0.9 decrease in f and causes a net increase in the appearance probability. We note that this type of situation does not arise when the decay function is exponential because changes in weight are linked to changes in f uniformly across all batches.

A general solution to this problem is to choose ρ_k as large as possible to maximize the sample size, while ensuring monotonicity in the appearance probabilities. Specifically, for a given i we want to choose ρ_k such that $\rho_k f(\alpha_{i,k}) \leq \rho_{k-1} f(\alpha_{i,k-1})$. The maximum feasible value is $\rho_{k,i}^* = \rho_{k-1} f(\alpha_{i,k-1})/f(\alpha_{i,k})$ as in our example. Because the inequality must hold for every $i \leq k$, the overall maximum feasible value is $\rho_1^* = 1$ and $\rho_k^* = \min_{i \leq k} \rho_{k,i}^*$ for $k > 1$. (The definition of ρ_k^* must actually be adjusted slightly to deal with the consolidated sample, as in line 15 of Algorithm 8.) As indicated by the previous example, ρ_k^* imposes an upper bound on the sample weight, so that direct use of ρ_k^* will produce samples whose expected size is generally less than $\min(W_k, n)$ at each t_k . In particular, the algorithm may return a sample of size less than n at a time t_k even when $W_k > n$. We deal with this problem by increasing the maximum sample weight, thereby trading off storage and sample size stability. In detail, we set $n' > n$ as the maximum sample weight, so that maximum sample weight and maximum sample size no longer coincide as in the exponential case. We then set

$$\rho_k = \min(1, n'/W_k, \rho_k^*) \quad (13)$$

for $k \geq 1$. If the total sample weight \mathbb{C}_k exceeds n at time t_k then we downsample to reduce the weight to n before outputting S_k (line 28). Lines 12–16 in Algorithm 8 implement the computation of ρ_k . Arguments essentially identical to those in Theorem 4.3 show that $\mathbb{C}_k = \rho_k W_k$ for all $k \geq 1$, where ρ_k is now defined as above. Since $\rho_k = \min(1, n'/W_k, \rho_k^*)$, we have that the total sample weight is bounded by $\min(W_k, n')$. As discussed in Section 4.3.2, we can bound the number of latent samples in \mathbb{L}_k , so that the overall sample footprint is bounded.

The following proposition helps in understanding the sample size dynamics and the precise way in which the choice of n' trades off storage for sample-size stability. In the proposition, we set $W_k^* = \max_{1 \leq i \leq k} W_i$.

PROPOSITION 4.7. *Algorithm 8 has the following properties:*

- (i) *If $W_k^* \leq n'$, then $\rho_i = 1$ for $1 \leq i \leq k$;*
- (ii) *If $W_k^* > n'$, then $\rho_k \geq n'/W_k^*$; and*
- (iii) *If the decay function f is strictly decreasing and $\rho_k = \rho_k^* < \min(1, n'/W_k)$, then $\rho_k > \rho_{k-1}$.*

Observe that, by (13), $\rho_k \leq 1$ for all k ; when $\rho_k = 1$ (so that $W_k \leq n'$), R-TBS achieves the largest possible sample weight of $\mathbb{C}_k = W_k$. Assertion (i) of the proposition implies that, as in the case of exponential decay, R-TBS returns the maximum possible sample size until the total weight first exceeds n' . As indicated by the previous example, one problematic scenario occurs when W_k exceeds n' and then drops rapidly but stays above n . After such an “adversarial” drop, the necessity of ejecting existing items and filtering out newly arrived items to enforce (1) means that R-TBS returns a sample of size smaller than n . In the case where $W_k^* > n'$, it follows from assertion (ii) that the UNIONED latent sample has a sample weight of $\mathbb{C}_k = \rho_k W_k \geq n' W_k / W_k^*$. This weight will exceed n —so that we can return a final sample of size n to the user after downsampling—as long as $W_k \geq (n/n') W_k^*$. Thus the larger the value of n' , the larger the adversarial drop from W_k^* that can be tolerated without reducing the returned sample size below its optimal value of n . As indicated by assertion (iii), we see that, when $\rho_k < \min(1, n'/W_k)$, so that $\rho_k = \rho_k^*$ and suboptimal sample sizes are being returned, the ρ_k values will increase toward the optimal values of $\min(1, n'/W_k)$ as time progresses, so that the algorithm will steadily “recover” from the adversarial drop until the next adversarial drop occurs. An analogous situation occurs when the total weight decreases to a value less than n at an adversarial drop. In this case the sample weight \mathbb{C}_k falls below the desired value of W_k but then recovers as time goes on.




Fig. 3. R-TBS sample size fluctuations over time with quadratic decay; $n = 1000$

This behavior is illustrated in Figure 3. In the figure, arriving batch sizes rise and fall periodically, over periods of length $m = 2000$ batches. In each period, the first $(2/3)m$ batches each contain b_1 items and the last $(1/3)m$ batches each contain b_2 items, where $b_1 < b_2$, leading to a sequence of adversarial drops. In Figures 3(a) and 3(b), we have $b_1 = 100$, $b_2 = 300$ and the total weight W_k exceeds $n = 1000$ at all times. Each sharp spike upwards in W_k (due to the sudden jump in batch size) causes ρ_k to fall, in order to maintain the proper inclusion probabilities. At each subsequent sharp decline in W_k , the sample size \mathbb{C}_k dips below n , even though the total weight exceeds n . The ρ_k values then slowly recover, approaching 1, until the next spike causes the cycle to repeat. Comparing Figures 3(a) and 3(b), we see that increasing the maximum sample weight from $n' = n$

to $n' = 2n$ largely ameliorates the dips in the sample size. The sample-size dynamics in Figure 3(c) are similar, but here $b_1 = 50$ and $b_2 = 250$, so that the total weight W_k occasionally falls below n . Just after each sharp decline in W_k , there is a small dip where $C_k < W_k < n$ for a small number of sequential values of k ; in the case of exponential decay we would have $C_k = W_k < n$ at these time points. Because we take $n' = 2n$, the magnitude of the dips is small.

5 DISTRIBUTED TBS ALGORITHMS

In this section, we describe how to implement distributed versions of T-TBS and R-TBS to handle large volumes of data.

5.1 Overview of Distributed Algorithms

The distributed T-TBS and R-TBS algorithms, denoted as D-T-TBS and D-R-TBS respectively, need to distribute large data sets across the cluster and parallelize the computation on them.

For exponential decay, at any point in time, all items have exactly the same decay rate for their appearance probabilities, regardless of age, so there is no need to keep track of an item's age in the sample. This nice "memoryless" property of the exponential function makes the distributed implementation of both algorithms easier. We first describe D-T-TBS and D-R-TBS for exponential decay, and then discuss the extensions for general decay functions in Section 5.4.

Overview of D-T-TBS: The implementation of the D-T-TBS algorithm is very similar to the simple distributed Bernoulli time-biased sampling algorithm in [35]. It is embarrassingly parallel, requiring no coordination. At each time point t_k , each worker in the cluster downsamples its partition of the sample with probability p , downsamples its partition of \mathcal{B}_k with probability q , and then takes a union of the resulting data sets.

Overview of D-R-TBS: This algorithm, unlike D-T-TBS, maintains a bounded sample, and hence is not embarrassingly parallel. D-R-TBS first needs to aggregate the local partition sizes for the incoming batch \mathcal{B}_k to compute the total batch size $|\mathcal{B}_k|$ and calculate the new total weight W_k . Then, based on $|\mathcal{B}_k|$, W_k , and the current sample weight C_k , D-R-TBS computes the downsample rate for the items in the reservoir, as well as the downsample rate for the items in \mathcal{B}_k . After that, D-R-TBS chooses the items in the reservoir to delete through a DOWNSAMPLE operation, selects items in \mathcal{B}_k (also via DOWNSAMPLE), inserts the selected items into the reservoir (via UNION), and finally generates the sample (via GETSAMPLE). Each of the expensive operations DOWNSAMPLE, UNION, and GETSAMPLE is performed in a distributed manner. They each require the master to coordinate among the workers. GETSAMPLE and UNION operations are relatively straightforward. The most challenging part of D-R-TBS lies in choosing items to delete from the reservoir and selecting new items to insert into the reservoir. In Section 5.3, we introduce two alternative approaches to determine the deleted and inserted items.

Both D-T-TBS and D-R-TBS periodically checkpoint the reservoir as well as other system state variables to ensure fault tolerance—the fault tolerance of the distributed algorithms in Spark is discussed in the Online Appendix. The implementation details for D-T-TBS are mostly subsumed by those for D-R-TBS, so we focus on the latter.

5.2 Distributed Data Structures

There are two important data structures in the D-R-TBS algorithm: the incoming batch and the reservoir. Conceptually, we view an incoming batch \mathcal{B}_k as an array of slots numbered from 1 through $|\mathcal{B}_k|$, and the reservoir as an array of slots numbered from 1 through $\lfloor C_k \rfloor$ containing full items plus a special slot for the partial item. For both data structures, data items need to be distributed into partitions due to the large data volumes. Therefore, the slot number of an item maps to a specific partition ID and a position inside the partition.




Fig. 4. Design choices for implementing the reservoir

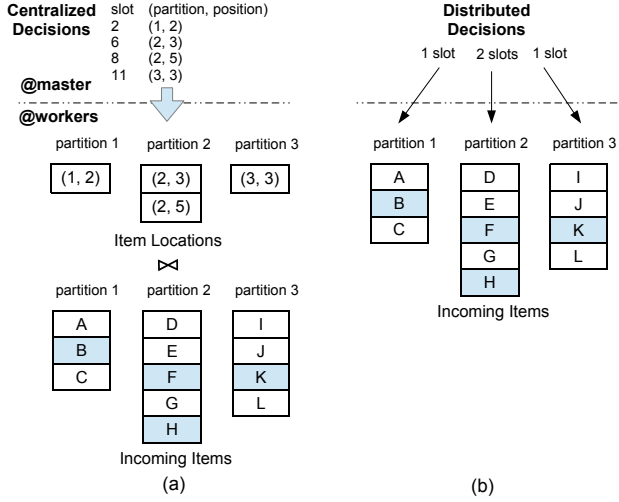


Fig. 5. Retrieving insert items

The incoming batch usually comes from a distributed streaming system, such as Spark Streaming; the actual data structure is specific to the streaming system, e.g., an incoming batch is stored as an RDD in Spark Streaming. As a result, the partitioning strategy of the incoming batch is opaque to the D-R-TBS algorithm. Unlike the incoming batch, which is read-only and discarded at the end of each time period, the reservoir data structure must be continually updated. An effective strategy for storing and operating on the reservoir is thus crucial for good performance. We now explore alternative approaches to implementing the reservoir.

Distributed in-memory key-value store: One quite natural approach implements the reservoir using an off-the-shelf distributed in-memory key-value store, such as Redis [31] or Memcached [27]. In this scheme, each item in the reservoir is stored as a key-value pair, with the slot number as the key and the item as the value. The partial item has a special slot number such as -1. Inserts and deletes to the reservoir naturally translate into put and delete operations to the key-value store.

There are three major limitations to this approach. First, the hash-based or range-based data-partitioning scheme used by a distributed key-value store yields reservoir partitions that do not correlate with the partitions of incoming batch. As illustrated in Figure 4(a), when items from a given partition of an incoming batch are inserted into the reservoir, the inserts touch many (if not all) partitions of the reservoir, incurring heavy network I/O. Second, key-value stores incur

unnecessary concurrency-control overhead. For each batch, D-R-TBS already carefully coordinates the deletes and inserts so that no two delete or insert operations access the same slots in the reservoir and there is no danger of write-write or read-write conflicts. Finally, the key-value store approach requires an explicit slot number for each item. As a result, D-R-TBS needs to take extra care to make sure that after deletes and inserts of reservoir items, the slot numbers are still unique and contiguous, e.g. by recycling the slot numbers of deleted items for new inserts. The burden of keeping track of delete and insert slot numbers falls on the master node.

Co-partitioned reservoir: An alternative approach implements a distributed in-memory data structure for the reservoir so as to ensure that the reservoir partitions coincide with the partitions from incoming batches, as shown in Figure 4(b). This can be achieved in spite of the unknown partitioning scheme of the streaming system. Specifically, the reservoir is initially empty, and all items in the reservoir are from the incoming batches. Therefore, if an item from a given partition of an incoming batch is always inserted into the corresponding “local” reservoir partition and deletes are also handled locally, then the co-partitioning and co-location of the reservoir and incoming batch partitions is automatic. In addition, if the incoming batch is evenly distributed across the different partitions (which is often the case in practice), then the co-partitioned reservoir is also evenly distributed. For our experiments, we implemented the co-partitioned reservoir in Spark using the in-place updating technique for RDDs in [35]; see the Online Appendix.

Note that, at any point in time, a given (conceptual) slot number in the reservoir maps to a specific partition ID and a position inside the partition. The mapping between a specific full item and its current slot number may change over time due to reservoir insertions and deletions. This does not cause any statistical issues, because the set-based R-TBS algorithm is oblivious to specific slot numbers. Thus the master only needs to keep track of the size of each partition and the position of the partial item. In Section 6.3, we experimentally compare the key-value store and the co-partitioned reservoir approaches.

5.3 Choosing Items to Delete and Insert

In order to bound the reservoir size, D-R-TBS requires careful coordination among workers when choosing the set of items to delete from, and insert into, the reservoir. At the same time, D-R-TBS must ensure the statistical correctness of random number generation and random permutation operations in the distributed environment. We consider two possible approaches.

Centralized decisions: In the most straightforward approach, the master makes centralized decisions about which items to delete and insert. For deletes, the master generates slot numbers of the items in the reservoir to be deleted, which are then mapped to the actual data locations in a manner that depends on the representation of the reservoir (key-value store or co-partitioned reservoir). For inserts, the master generates the slot numbers of the incoming items \mathcal{B}_k at time t_k that need to be inserted into the reservoir. Suppose that \mathcal{B}_k comprises $m \geq 1$ partitions. Each generated slot number $i \in \{1, 2, \dots, |\mathcal{B}_k|\}$ is mapped to a partition p_i of \mathcal{B}_k (where $1 \leq p_i \leq m$) and a position r_i inside partition p_i . Denote by Q the set of “item locations”, i.e., the set of (p_i, r_i) pairs. In order to perform the inserts, D-R-TBS needs to first retrieve the actual items based on the item locations. This can be achieved with a join-like operation between Q and \mathcal{B}_k , with the (p_i, r_i) pair matching the actual location of an item inside \mathcal{B}_k . To optimize this operation, we make Q a distributed data structure and use a customized partitioner to ensure that all pairs (p_i, r_i) with $p_i = j$ are co-located with partition j of \mathcal{B}_k for $j = 1, 2, \dots, m$. Then a co-partitioned and co-located join can be carried out between Q and \mathcal{B}_k , as illustrated in Figure 5(a) for $m = 3$. The resulting set of retrieved insert items, denoted as S , is also co-partitioned with \mathcal{B}_k as a by-product. After that, the actual deletes and inserts are then carried out depending on how reservoir is stored, as discussed below.

When the reservoir is implemented as a key-value store, the deletes can be directly applied based on the slot numbers. For inserts, the master takes each generated slot number of an item in \mathcal{B}_k and chooses a companion destination slot number in the reservoir into which the \mathcal{B}_k item will be inserted. This destination reservoir slot might currently be empty due to an earlier deletion, or might contain an item that will now be replaced by the newly inserted batch item. After the actual items to insert are retrieved as described previously, the destination slot numbers are used to put the items into the correct locations in the key-value store.

When the co-partitioned reservoir is used, the delete slot numbers in the reservoir are mapped to (p_i, r_i) pairs of partitions of the reservoir and positions inside the partitions. As with inserts, we again use a customized partitioner for the set of pairs \mathcal{R} such that deletes are co-located with the corresponding reservoir partitions. Then a join-like operation on \mathcal{R} and the reservoir performs the actual delete operations on the reservoir. For inserts, we simply use another join-like operation on the set of retrieved insert items \mathcal{S} and the reservoir to add the corresponding insert items to the co-located partition of the reservoir. In this approach, we don't need the master to generate destination reservoir slot numbers for these insert items, because we view the reservoir as a set when using a co-partitioned reservoir data structure.

Distributed decisions: The above approach requires the master to generate a large number of slot numbers, so we now explore an alternative approach that offloads the slot number generation to the workers while still ensuring the statistical correctness of the computation. This approach has the master choose only the number of deletes and inserts per worker according to an appropriate multivariate hypergeometric distribution. For deletes, each worker chooses random victims from its local partition of the reservoir based on the number of deletes given by the master. For inserts, the worker receives the number of inserts I and then randomly and uniformly selects I items from its local partition of the incoming batch \mathcal{B}_k . Figure 5(b) depicts how the insert items are retrieved under this decentralized approach. We use the technique in [17] for parallel pseudo-random number generation.

The foregoing distributed decision making approach works only when the co-partitioned reservoir data structure is used. This is because the key-value store representation of the reservoir requires a target reservoir slot number for each insert item from the incoming batch, and the target slot numbers have to be generated in such a way as to ensure that, after the deletes and inserts, all of the slot numbers are still unique and contiguous in the new reservoir. This requires a lot of coordination among the workers, which inhibits truly distributed decision making.

5.4 Extensions to Generalized TBS Algorithms

We now discuss the extensions to the above distributed implementations to make them work for general (non-exponential) decay functions.

Changes to the reservoir data structure: Use of general decay functions requires some significant changes to the reservoir data structure. In the absence of the special memoryless property for exponential decay functions, the reservoir needs to keep track of the age for each item, since the decay rate for an item depends explicitly on its age. In addition, the generalized D-R-TBS algorithm also requires storage of multiple latent samples, and hence multiple partial items, in the reservoir—one per each age, up to the cutoff age, plus one for the consolidated latent sample of older items.

In the key-value store approach, D-R-TBS needs to record the arrival time of each item in the sample by augmenting the value component of the key-value pair that represents the item; the item's age can then be calculated on the fly. D-R-TBS also needs to maintain multiple special key-value pairs for the partial items. As discussed in Section 5.3, the key-value store approach only works when centralized decisions determine which items to delete from, and insert into,

the reservoir. Thus for general decay functions, the master node needs to track the arrival time corresponding to each slot number of the reservoir, so that it can apply different decay rates to differently aged items when deleting items from the reservoir.

In the co-partitioned reservoir approach, D-R-TBS also needs to record the item arrival times and to support multiple latent samples corresponding to different ages. In the implementation, we can either simply add an arrival-time field to each item in the reservoir or organize the items for each arrival time together in each partition. We choose the latter for the ease of the downsampling process. Assuming that each incoming batch is evenly distributed across partitions, then the items corresponding to each arrival time are also evenly distributed across the reservoir partitions, and hence the overall reservoir structure is well balanced. Finally, the master needs to record the number of items for each arrival time in each partition of the reservoir and the positions of the partial items for each arrival time. This overhead is much smaller than the overhead of tracking the arrival time corresponding to each slot number, as needed in the key-value-store approach. Thus, for general decay functions, the co-partitioned reservoir approach dominates the key-value store approach even more than in the exponential decay setting.

Changes to item insertion and deletion: The presence of general decay functions adds complexity to the process of choosing items to delete from, and insert into, the reservoir, since different decay rates apply to different latent samples in the reservoir. As mentioned before, the centralized-decision approach requires the master to record the arrival time corresponding to each slot number of the reservoir in order to select items to delete according to the correct probabilities. Moreover, as in the exponential case, special care needs to be taken to make sure that the slot numbers remain unique and contiguous after deletes and inserts. In contrast, for the distributed-decision approach, the master merely needs to decide the number of items to delete for each arrival time in each partition; the workers carry out the actual deletes. The remaining aspects of the distributed implementation stay the same.

6 EXPERIMENTS

In this section, we study the empirical performance of distributed implementations of the R-TBS and T-TBS algorithms, and demonstrate the potential benefit of using them for model retraining in online model management.

Experimental Setup: We implemented R-TBS and T-TBS on Spark (the Online Appendix contains Spark-specific implementation details). All performance experiments were conducted on a cluster of 9 ProLiant DL160 G6 servers. Each has two twelve-core Intel Xeon X5650 CPUs at 2.66GHz, with 15GB of RAM and a single 7200 RPM 500GB hard drive. Servers are interconnected using a 1 Gbit Ethernet and each server runs CentOS release 6.5, Java 1.8 and Spark 2.3. One server is dedicated to run the Spark coordinator and, each of the remaining servers runs a single Spark worker with parallelism 10, along with 10 GB of dedicated memory. All other Spark parameters are set to their default values. We used Memcached 1.4.4 as the key-value store in our experiments.

We note that the experimental setup in this paper differs from that in our previous work [18]. Besides using a different hardware configuration, we also upgraded Spark from version 1.6 to 2.3 and Memcached from version 1.4.33 to 1.4.4. All reported experimental results correspond to the current hardware and software configuration. The performance of Spark 2.3 has dramatically improved over 1.6, and all algorithms benefited from these improvements. Although the actual runtimes of each algorithm changed, the qualitative results in the exponential case are generally consistent with those in [18]. Interestingly, we observed a reduction in the performance gap between R-TBS and T-TBS; see below.

For all experiments, data was streamed in from HDFS using Spark Streaming's microbatches. Decay occurs according to a time scale such that the batch-arrival interval is $\Delta = 1$ in the decay

formulas. We report run time per batch as the average over 100 batches, discarding the first round from this average because of Spark startup costs. We experiment with two versions of R-TBS. The first, *R-TBS exp*, refers to the specialized version of R-TBS for exponential decay with $f(\alpha) = e^{-\lambda\alpha}$. The second, *R-TBS poly*, refers to the generalized R-TBS with a shifted polynomial decay function $f^{(d)}(\alpha) = (1+d)^s/(1+d+\alpha)^s$; as discussed in Section 3.1, the shift is necessary so that item's weights do not become too small too quickly. We similarly experiment with T-TBS exp and T-TBS poly. Unless otherwise stated, default values of $\lambda = 0.07$ and $n = 2 \times 10^7$ are used for R-TBS exp. For R-TBS poly, we use $n = 2 \times 10^7$ unless otherwise stated and the other parameters take on values $(s, d, n', n, \lambda, \delta_1, \delta_2) = (2, 10, 2n, n, 0.1, 0.01, 0.001n)$. Finally, for the runtime experiments, each batch contains 10 million items unless otherwise stated.

6.1 Summary of Results

Before we dive into the detailed experiments, we first highlight some of the major takeaways from our empirical studies.

Sample Size Behavior: Section 6.2 first compares the sample size behavior of T-TBS and R-TBS. The results empirically validate our previous assertions about T-TBS: although it is a much simpler algorithm than R-TBS, it suffers from sample overflow and/or underflow, especially when the mean batch size changes over time or the batch size fluctuates strongly. R-TBS maintains a bounded, relatively stable sample size throughout.

Performance of Distributed Implementation: Section 6.3 evaluates the various distributed implementation strategies described in Section 5. The most highly optimized implementation exhibits an almost 10x performance benefit relative to the most naive implementation.

ML Applications: Sections 6.4 through 6.6 compare the accuracy and robustness of R-TBS, simple sliding windows (SW), and uniform sampling (Unif) in three representative ML applications: a kNN classifier, a naive Bayes classifier, and a linear regression predictor. These applications span both parametric and nonparametric approaches, as well as both classification and prediction tasks. We find that R-TBS tends to have better accuracy and robustness than SW and Unif in the presence of reoccurring patterns in both single-change and periodic-change regimes. In addition, for linear regression, we compared retraining of models using R-TBS with online learning approaches adapted to batched streaming inputs. These online approaches, which work only for parametric models, can adapt to drastic changes more quickly than R-TBS, yielding better accuracy, but when either the changes are not as drastic or the ML model is complex, with a large number of parameters, R-TBS performs better.




Fig. 6. Sample size behavior for T-TBS and R-TBS under exponential decay; λ = decay rate and ϕ = batch size multiplier




Fig. 7. Sample size behavior for T-TBS and R-TBS under quadratic decay; ϕ = batch size multiplier

6.2 Sample Size Behavior

We first compare the sample size behavior of T-TBS and R-TBS under a variety of batch size regimes. Throughout, we write $B_k = |\mathcal{B}_k|$. Figures 6 and 7 show sample size behavior over time for the two algorithms under exponential and quadratic decay. The key challenge to T-TBS is that the value of the mean batch size b must be specified in advance; consequently, the algorithm cannot handle dynamic changes in b without losing control of either the decay rate or the sample size.

In Figure 6(a), for example, the (deterministic) batch size is initially fixed and the algorithm is tuned to a target sample size of 1000, with a decay rate of $\lambda = 0.05$. At $k = 200$, the batch size starts to increase (with $B_{k+1} = \phi B_k$ where $\phi = 1.002$), leading to an overflowing sample, whereas R-TBS maintains a constant sample size.

Even in a stable batch-size regime with constant batch sizes (or, more generally, small variations in batch size), R-TBS can maintain a constant sample size whereas the sample size under T-TBS fluctuates in accordance with Theorems 3.1 and 3.2; see Figure 6(b) for the case of a constant batch size $B_k \equiv 100$ with $\lambda = 0.1$.

Large variations in the batch size lead to large fluctuations in the sample size for T-TBS; in this case the sample size for R-TBS is bounded above by design, but large drops in the batch size can cause drops in the sample size for both algorithms; see Figure 6(c) for the case of $\lambda = 0.1$ and i.i.d. uniformly distributed batch sizes on $[0, 200]$ so that $E[B_k] \equiv 100$. Similarly, as shown in Figure 6(d), systematically decreasing batch sizes will cause the sample size to shrink for both T-TBS and R-TBS. Here, $\lambda = 0.01$ and, as with Figure 6(a), the batch size is initially fixed and then starts to change at $k = 200$, with $\phi = 0.8$ in this case. This experiment—and others, not reported here, with varying values of λ and ϕ —indicate that R-TBS is more robust to sample underflows than T-TBS.

The results for polynomial decay in Figure 7 are similar to those for exponential decay. In the shrinking-batch-size scenario, note that R-TBS has a harder time maintaining a full sample under quadratic decay than under exponential decay (but still does better than T-TBS). The lower sample sizes are a direct consequence of the fact that different items decay at different rates, as discussed in detail in Section 4.3.3.

So T-TBS has much more trouble maintaining a target sample size than R-TBS, especially when batch sizes fluctuate unpredictably. As pointed out in Section 3.2, however, when T-TBS is applicable, it is much simpler and faster than R-TBS.

6.3 Runtime Performance

Comparison of TBS Implementations: Figure 8 shows the average runtime per batch for four implementations of distributed R-TBS exp with different design choices. The design decisions discussed are whether to use centralized or distributed decisions in choosing items to insert and delete (abbreviated as "Cent" and "Dist", respectively), whether to implement the reservoir using a




Fig. 8. Per-batch runtime comparison of different implementations of R-TBS with exponential decay




Fig. 9. Per-batch runtime of R-TBS and T-TBS with exponential decay and polynomial decay




Fig. 10. Scale out of D-R-TBS




Fig. 11. Scale up of D-R-TBS




Fig. 12. Effect of Changing Decay Parameter

key-value store or a co-partitioned reservoir scheme (abbreviated as "KV" and "CP"), and whether to subsample the incoming batch using the standard repartition join or using a copartitioned join (abbreviated as "RJ" and "CJ") under centralized decision scheme. These design decisions are discussed in more detail in Section 5.

The first two implementations in Figure 8 both use the key-value store representation for the reservoir together with the centralized decision strategy for determining inserts and deletes. They only differ in how the inserted items are retrieved when subsampling the incoming batch. The first uses the repartition join, whereas the second uses the customized partitioner and co-located join. This optimization effectively cuts the runtime in half, but the KV representation of the reservoir still requires the inserted items to be written across the network to their corresponding reservoir location. The third implementation employs the co-partitioned reservoir instead, resulting in a speedup of 2.75x. The fourth implementation further employs the distributed decision for choosing items to delete and insert. This yields a further 1.46x speedup. The combination of co-partitioned reservoir scheme and the distributed decision making for inserting and deleting items always yields the best performance for R-TBS and T-TBS with different decay functions, so we use this combination for the remaining experiments.

In Figure 9, we show the per batch runtimes for R-TBS and T-TBS with both exponential decay and polynomial decay. R-TBS exp and R-TBS poly have very similar runtime performance, with R-TBS exp being slightly faster; a similar observation holds for T-TBS. Since T-TBS is embarrassingly parallelizable, it is faster than R-TBS—though, as mentioned previously, the relative performance advantage of T-TBS is smaller than the result reported in [18] due to improvements in Spark. In any case, as discussed in Section 3, T-TBS is faster, but only works under a very strong restriction on the data arrival rate, and can suffer from occasional memory overflows; see Figures 6 and 7. In

contrast, R-TBS has more robust sample-size behavior and works in realistic scenarios where it is hard to predict the data arrival rate.

Scalability of R-TBS: Figure 10 shows how R-TBS with exponential and non-exponential (quadratic in this case) decay functions scale with the number of workers. The two implementations have very similar performance. Initially, both versions of R-TBS scale out very nicely and see linear speedup from an increase in workers. However, beyond 4 workers, the marginal benefit from additional workers is small, because the coordination and communication overheads, as well as the inherent Spark overhead, become prominent. For the same reasons, in the scale-up experiment in Figure 11, both runtimes stay roughly constant until the batch size reaches 1 million items and then increase sharply at 10 and 100 million items. This is because processing the streaming input and maintaining the sample start to dominate the coordination and communication overhead. With 8 workers, our implementation of R-TBS can handle a data flow comprising 100 million items arriving approximately every 16 seconds.

R-TBS Runtime Robustness: Figure 12 shows the impact of changes in the exponential decay parameter on runtime; as can be seen, the impact is negligible. Similar results hold for changes in the exponent for R-TBS poly (quadratic, cubic etc.). These results might seem counterintuitive, since changes in the decay parameter have a substantial impact on the number of items inserted into and deleted from the reservoir. However, in our optimized distributed implementation of R-TBS (with co-partitioned reservoir and distributed decision making), these are not expensive operations. All inserts and deletes happen locally and only affect local memory. In comparison, the cost of reading the incoming batch across the network or from disks, as well as the communication overhead between the Spark master and Spark workers, are much more expensive. Overall, the processing cost is dominated by the cost of reading in incoming batches, which is linear in the batch size, and thus the average runtime depends only on the expected batch size, regardless of the batch size variability. While not shown here, similar results hold for skew in the number of items per incoming batch.

6.4 Application: Classification using kNN

Our first ML model is a kNN classifier, where a class is predicted for each item in an incoming batch by taking a majority vote of the classes of the k nearest neighbors in the current sample, based on Euclidean distance; the sample is then updated using the batch. To generate training data, we first generate 100 class centroids uniformly in a $[0, 80] \times [0, 80]$ rectangle. Each data item is then generated from a Gaussian mixture model and falls into one of the 100 classes. Over time, the data generation process operates in one of two “modes”. In the “normal” mode, the frequency of items from any of the first 50 classes is five times higher than that of items in any of the second 50 classes. In the “abnormal” mode, the frequencies are five times lower. Thus the frequent and infrequent classes switch roles at a mode change. We generate each data point by randomly choosing a ground-truth class c_i with centroid (x_i, y_i) according to relative frequencies that depend upon the current mode, and then generating the data point’s (x, y) coordinates independently as samples from $N(x_i, 1)$ and $N(y_i, 1)$. Here $N(\mu, \sigma)$ denotes the normal distribution with mean μ and standard deviation σ .

In this experiment, the batch sizes are deterministic with $b = 100$ items. The sample size for both R-TBS and uniform reservoir sampling (Unif) is 1000, and the sliding window (SW) contains the last 1000 items; thus all methods use the same amount of data for retraining. The hyperparameter k is tuned individually for each of R-TBS, SW, and Unif sampling schemes; we choose the value that minimizes the misclassification percentage. In each run, the sample is warmed up by processing 100 normal-mode batches before the classification task begins. We test R-TBS with both exponential and shifted-polynomial decay functions as described in Section 6.3. We use

$\lambda = 0.07$ for R-TBS exp, unless otherwise stated. For R-TBS poly, we use the parameter values of $(s, d, n', n, \lambda, \delta_1, \delta_2) = (2, 10, 2000, 1000, 0.1, 0.01, 1)$. We focus on two types of temporal patterns in the data, as described below.

Single change: Here we model the occurrence of a singular event. The data is generated in normal mode up to $k = 10$ (time is measured here in number of units after warm-up), then switches to abnormal mode, and finally at $t = 20$ switches back to normal. As can be seen in Figure 13(a), the misclassification rate (percentage of incorrect classifications) for all sampling schemes increases from around 15% to roughly 40% when the distribution becomes abnormal. Both versions of R-TBS as well as SW adapt to the change, recovering to around 13% misclassification rate after $t = 16$, with R-TBS poly and SW adapting slightly better than R-TBS exp. In comparison, Unif does not adapt at all. But, when the distribution snaps back to normal, the error rate of SW rises sharply to 40% before gradually recovering, whereas error rates of R-TBS exp and R-TBS poly stay low around 13% and 15% throughout. These results show that R-TBS is indeed more robust: although sometimes slightly more sluggish than SW in adapting to changes, R-TBS avoids wild fluctuations in classification error as with SW.

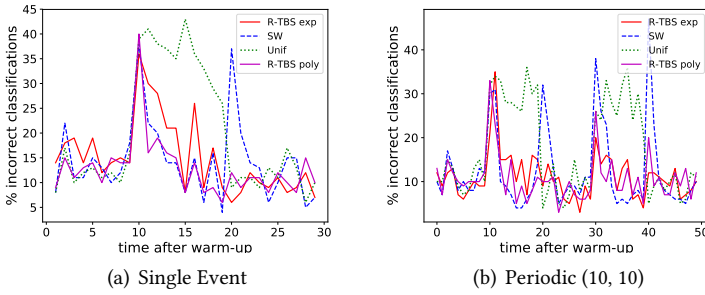


Fig. 13. Misclassification rate (percent) for kNN




Fig. 14. Misclassification rate (percent) for kNN

Periodic change: For this temporal pattern, the changes from normal to abnormal mode are periodic, with δ normal batches alternating with η abnormal batches, denoted as $\text{Periodic}(\delta, \eta)$, or $P(\delta, \eta)$ for short. Figure 13(b) shows the misclassification rate for $\text{Periodic}(10, 10)$. Experiments on other periodic patterns demonstrate similar results; see Figures 14(a) and 14(b). The robust

behavior of R-TBS described above manifests itself even more clearly in the periodic setting. Note, for example, how both R-TBS versions react significantly better to the renewed appearances of the abnormal mode. Observe that the first 30 batches of Periodic(10, 10) display the same behavior as in the single event experiment in Figure 13(a). We therefore focus primarily on the Periodic(10, 10) temporal pattern for the remaining experiments.

Robustness and Effect of Decay Parameter: In the context of online model management, we need a sampling scheme that delivers high overall prediction accuracy and, perhaps even more importantly, robust prediction performance over time. Large fluctuations in the accuracy can pose significant risks in applications, e.g., in critical IoT applications in the medical domain such as monitoring glucose levels for predicting hyperglycemia events. To assess the robustness of the performance results across different sampling schemes, we use a standard risk measure called *expected shortfall (ES)* [26, p. 70]. ES measures downside risk, focusing on worst-case scenarios. Specifically, the $z\%$ ES is the average value of the worst $z\%$ of cases.

For each of 30 runs and for each sampling scheme, we compute the 10% ES of the misclassification rate (expressed as a percentage) starting from $t = 20$, since all three sampling schemes perform poorly (as would be expected) during the first mode change, which finishes at $t = 20$. Table 1 lists both the *accuracy*, measured in terms of the average misclassification rate, and the *robustness*, measured as the average 10% ES, of the kNN classifier over 30 runs across different temporal patterns. To demonstrate the effect of the exponential decay parameter λ and polynomial decay parameter s on model performance, we also include numbers for different λ and s values in Table 1. Additionally, while not reported for reasons of space, we varied the shift parameter d for polynomial decay and found similar results to those below for all values of d between 6 and 14.

In terms of accuracy, Unif is always the worst by a large margin. R-TBS and SW have similar accuracies, with R-TBS being slightly more accurate. On the other hand, for robustness, SW is almost always the worst, with ES ranging from 1.4x to 2.5x the ES of R-TBS. Mostly, Unif is also significantly worse than R-TBS, with ES ratios ranging from 1.3x to 1.7x. The only exception is the single-event pattern: since the data remains in normal mode after the abnormal period, time biasing becomes unimportant and Unif performs well. In general, R-TBS provides both better accuracy and robustness, with the best sampling schemes in terms of either accuracy or ES being R-TBS in all cases. Furthermore, this edge in accuracy and robustness is fairly stable across a wide range of λ and s values. Finally, when comparing R-TBS exp against R-TBS poly, we see that polynomial decay is slightly more accurate, but exponential decay has a slight edge with respect to robustness. Additionally, for both R-TBS exp and R-TBS poly, we find that the optimal decay value for ES is smaller than the optimal value for expected miss %, as might be expected, since a lower value leads to greater retention of older items. Overall, R-TBS exp and R-TBS poly yield comparable performance in ML models under dynamic data, but which is better depends on the experimental setup and so the choice of decay function will generally need to be driven by application requirements.

Varying batch size: We now examine model quality when the batch sizes are no longer constant. Overall, the results look similar to those for constant batch size. For example, Figure 15(a) shows results for a Uniform(0, 200) batch-size distribution, and Figure 15(b) shows results for a deterministic batch size that grows at a rate of 2% after warm-up. In both experiments, $\lambda = 0.07$ for R-TBS exp, $s = 2$ for R-TBS poly, and the data pattern is Periodic(10, 10). These figures demonstrate the robust performance of R-TBS in the presence of varying data arrival rates. Similarly, the average accuracy and robustness over 30 runs resembles the results in Table 1. For example, pick R-TBS exp with $\lambda = 0.07$ and the Periodic(10, 10) pattern. Then, the misclassification rate under uniform/growing batch sizes is 1.17x/1.14x that of R-TBS for SW, and 1.39x/1.38x for Unif. In addition, the ES is 1.68x/1.95x that of R-TBS for SW, and 1.64x/1.60x for Unif.

For the experiments in this section, we see that, in terms of model accuracy and robustness, both R-TBS exp and R-TBS poly perform similarly. We found that this observation holds generally, and so we will focus on R-TBS exp from now on.

Table 1. Accuracy and robustness of kNN performance

<i>Alg.</i>	Single Event		P(10,10)		P(20,10)		P(30,10)	
	Miss%	ES	Miss%	ES	Miss%	ES	Miss%	ES
R-TBS exp: $\lambda = 0.05$	17.1	16.8	16.1	22.1	15.3	24.4	15.1	25.9
R-TBS exp: $\lambda = 0.07$	16.5	17.3	15.3	21.3	14.9	24.0	14.4	25.2
R-TBS exp: $\lambda = 0.10$	15.7	18.5	15.1	22.1	14.7	24.9	14.7	26.9
R-TBS poly: $s = 1.8$	15.0	17.9	13.9	23.2	14.5	24.1	13.8	24.0
R-TBS poly: $s = 2.0$	14.9	18.9	14.2	23.7	14.1	24.2	13.7	24.6
R-TBS poly: $s = 2.2$	17.6	21.7	14.5	24.1	14.3	24.2	14.2	26.8
SW	19.2	42.1	17.1	41.7	16.1	39.8	15.9	38.3
Unif	21.3	18.3	25.4	34.8	19.6	35.7	19.0	35.8

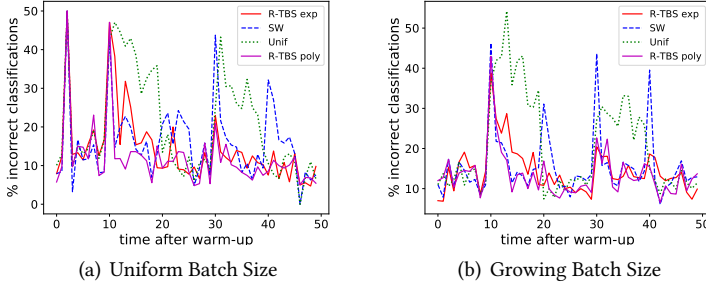


Fig. 15. Varying batch sizes for kNN classifier

6.5 Application: Linear Regression

We now assess the effectiveness of R-TBS for retraining regression models. The experimental setup is similar to kNN, with data generated in “normal” and “abnormal” modes. In both modes, data items are generated from the standard linear regression model $y = b_1x_1 + b_2x_2 + \epsilon$, with the noise term ϵ distributed according to a $N(0, 1)$ distribution. In normal mode, $(b_1, b_2) = (4.2, -0.4)$ and in abnormal mode, $(b_1, b_2) = (-3.6, 3.8)$. In both modes, x_1 and x_2 are generated according to a Uniform(0, 1) distribution. As before, the experiment starts with a warm-up of 100 “normal” mode batches and each batch contains 100 items. For each sampling scheme, the model is trained using a ridge-regression penalty hyperparameter that minimizes mean-squared error.

Saturated samples: Figure 16(a) shows the performance of R-TBS, SW, and Unif under the Periodic(10, 10) pattern with a maximum sample size of 1000 for each technique, and $\lambda = 0.10$ for R-TBS. We note that, for this sample size and temporal pattern, the R-TBS sample is always saturated. (This is also true for all of the prior experiments.) The results echo that of the previous section, with R-TBS exhibiting slightly better prediction accuracy on average, and significantly better robustness, than the other methods. The mean square errors (MSEs) across all data points for R-TBS, Unif, and SW are 3.28, 4.41, 3.98 respectively, and their 10% ES of the MSEs are 8.16, 11.17, 10.65 respectively.

Unsaturated Samples: We now investigate the case of unsaturated samples for R-TBS. We increase the target sample size to $n = 1600$. With a constant batch size of 100, and a decay rate




Fig. 16. Mean square error for linear regression




Fig. 17. Online methods react faster to changes but perform poorly as models get more complex

$\lambda = 0.10$, the reservoir of R-TBS is never full, stabilizing at 1051 items, whereas Unif and SW both have a full sample of 1600 items.

For the Periodic(10, 10) pattern, shown in Figure 16(b), SW has a window size large enough to keep some data from older time periods (up to 16 batches ago), making SW's robustness comparable to R-TBS (ES of 5.86 for SW and 6.01 for R-TBS). However, this amalgamation of old data also hurts its overall accuracy, with MSE rising to 4.16, as opposed to 3.43 for R-TBS. In comparison, the shape of R-TBS remains almost unchanged from Figure 16(a), and Unif behaves as poorly as before. When the pattern changes to Periodic(16, 16) as shown in Figure 16(c), SW doesn't contain enough old data, making its prediction performance suffer from huge fluctuations again, and the superiority of R-TBS is more prominent. In both cases, R-TBS provides the best overall performance, despite having a smaller sample size. This backs up our earlier claim that more data is not always better. A smaller but more balanced sample with good ratios of old and new data can yield better prediction performance than a large but unbalanced sample.

Comparison to Online ML Approaches: For parametric models, online approaches provide an alternative to data sampling in order to periodically adapt the model without retraining on all prior data. We consider two techniques: online gradient descent (OGD), where the current batch is used in a single step of *mini-batch gradient descent*, and proximal gradient methods (Prox), where a model is retrained on the current batch, but with a penalty term on the distance between the parameter values for the new and previous models; the idea is to prevent the model from changing too drastically at any given time step. We adapt both algorithms to batched streaming input, i.e., models are updated based on all the items in an incoming batch, instead of one item at a time.

As a first experiment, we use the same regression setup as before. We individually tune hyperparameters for each of the three models—the ridge regression penalty for R-TBS, the learning rate for OGD, and the proximal distance parameter for the proximal gradient approach—to minimize average L_2 error. Figure 17(a) shows a sample result. Prox performs the best, as it learns the new pattern in just a single round. In comparison, R-TBS and OGD learn the new pattern a bit more slowly. The average mean squared error for R-TBS, Prox, and OGD are 3.28, 2.04, and 3.05, respectively.

Our second experiment incorporates two changes. First, we scale down the regression coefficients b for each mode by a factor of $1/3$, which has the effect of reducing the jump in error at a mode change from $10\times$ normal error to only $2\times$. Additionally, we add another 198 parameters to the regression model, thus making the model harder to learn. As seen in Figure 17(b), both R-TBS and Prox still struggle at mode changes, but in non-change rounds R-TBS doesn't vary as wildly as Prox. OGD, not shown, performs so poorly that we didn't include it in the graph. This is due to the fact that it is hard to learn a pattern with many parameters in just a single pass over the data. The MSE values for this experiment are 1.39, 1.54, and 6.82 for R-TBS, Prox, and OGD, respectively. Figure 17(c) shows the effect of increasing the number of stable parameters in the model, with Prox performing better with fewer parameters and R-TBS performing better with a larger number of parameters. OGD again performs so poorly that it is left off the graph. Overall, from these experiments, along with others not shown, we observe that online models are more sensitive and react faster to drastic changes, but a data sampling approach based on R-TBS is more robust for milder changes and complex models with a large number of parameters.




Fig. 18. Misclassification rate (percent) for Naive Bayes

6.6 Application: Naive Bayes Classifier

In our final experiment, we evaluate the performance of R-TBS for retraining Naive Bayes models with the Usenet2 dataset (mlkd.csd.auth.gr/concept_drift.html), which was used in [20] to study classifiers coping with recurring contexts in data streams. This dataset contains a stream of 1500 messages on different topics from the 20 News Groups Collections [24]. They are sequentially presented to a simulated user who marks whether a message is interesting or not. The user's interest changes after every 300 messages. More details of the dataset can be found in [20].

Following [20], we use Naive Bayes with a bag-of-words model, and set the optimal parameters for SW with maximum sample size of 300 and batch size of 50. Since this dataset is rather small and contexts change frequently, we use the optimal value of 0.35 for λ . We find through experiments that R-TBS displays higher prediction accuracy for all λ in the range of $[0.1, 0.5]$, so precise tuning of λ is not critical. In addition, there is not enough data to warm up the models on different sampling schemes, so we report the model performance on all 30 batches. Similarly, we report 20% ES for this dataset, due to the limited number of batches. For each sampling scheme, the smoothing parameter that leads to the best misclassification rate is used.

The results are shown in Figure 18. The misclassification rate for R-TBS, SW, and Unif are 26.2%, 28.7%, and 29.1% and the 20% ES values are 43.2%, 48.8%, and 42.5%. Importantly, for this dataset the changes in the underlying data patterns are less pronounced than in the previous two experiments. Despite this, SW fluctuates wildly, yielding inferior accuracy and robustness. In contrast, Unif barely reacts to the context changes. As a result, Unif is very slightly better than R-TBS with respect to robustness, but at the price of lower overall accuracy. Thus, R-TBS is generally more accurate under mild fluctuations in data patterns, and its superior robustness properties manifest themselves as the changes become more pronounced.

7 RELATED WORK

Time-decay and sampling: Work on sampling with unequal probabilities goes back to at least Lahiri's 1951 paper [23]; see [29, Section 4] for some additional discussion of early work. A growing interest in streaming scenarios with weighted and decaying items began in the mid-2000's, with most of that work focused on computing specific aggregates from such streams, such as heavy-hitters, subset sums, and quantiles; see, e.g., [2, 9, 10]. The first papers on time-biased reservoir sampling with exponential decay are due to Aggarwal [1] and Efraimidis and Spirakis [14]; batch arrivals are not considered in these works. As discussed in Section 1, the sampling schemes in [1] are tied to item sequence numbers rather than the wall clock times on which we focus; the latter are more natural when dealing with time-varying data arrival rates.

Forward decay: Cormode et al. [11] consider an alternative framework for temporally-biased sampling based on an increasing "forward decay" function g , where the appearance probability at time t_k of an item arriving at time $t_i \leq t_k$ is proportional to $g(t_i)/g(t_k)$. This approach can be used to convert any recursively-defined weighted sampling scheme for a finite population to a streaming algorithm. Note that the decay behavior is quite different between forward and backward schemes. For example, a decay function $f(\alpha) = 1/(1 + \alpha)^2$ would naturally lead to forward decay function $g(t) = (1 + t)^2$. Observe, however, that in the backward scheme with uniform item arrival times and $\Delta = 1$, an item arriving at time t_i decays by a factor of $f(0)/f(1) = 1/4$ one time unit after it arrives, whereas in the forwards scheme, the decay factor is $g(i)/g(i + 1) = i^2/(i + 1)^2$, which becomes close to 1 for large i . As discussed in [11], forward decay schemes support a notion of "relative decay" where the weight of an item is determined by its fractional distance between an initial "landmark" time and the current time. In this paper we focus on backward decay because this latter notion is widely used and, we believe, easier for users to understand. Moreover, as indicated above, items decay relatively slowly in the forward scheme, which can cause ML algorithms to adapt too slowly to changes in the data. Finally, having experimented with forward decay versions of T-TBS and R-TBS, we observed that another issue with the forward decay in our setting is numerical stability. The work in [11] focused on analytical queries that are defined for items arriving during a specified, not-too-long time interval. In this setting, one can choose the forward-decay landmark time to be the time when the first item of interest arrives. For a sequence of such queries, the landmark time can be repeatedly shifted forward so that none of the item timestamps (measured relative to the landmark time) becomes too large. In our setting, our sample can retain items that have arrived arbitrarily long ago, so the only feasible landmark time is $t = 0$. This means that as time passes, the (absolute) time stamps become very large. The large timestamps pose numerical difficulties that are not amenable to normalization strategies. So again, we are motivated to focus on backward decay.

In the special case of an exponential decay function, the notions of backward and forward decay coincide, and so algorithms based on forward decay fall within our framework. (It is not hard to show that exponential decay functions are the only functions for which these notions coincide.) In this exponential setting, the authors in [11] provide a time biased reservoir sampling algorithm based on the A-Res weighted sampling scheme proposed in [14]. Rather than enforcing (1), however,

the algorithm enforces the (different) A-Res biasing scheme. In more detail, if s_i denotes the element at slot i in the reservoir, then the algorithm in [14] implements a scheme where an item x is chosen from a set of m elements to be at slot $i + 1$ in the reservoir with probability $w_x / (\sum_{j=1}^m w_j - \sum_{j=1}^i w_{s_j})$. From the form of this equation, it becomes clear that the resulting sampling algorithm violates (1). Indeed, Efraimidis [13] gives some numerical examples illustrating this point (in his comparison of the A-Res and A-Chao algorithms). We would argue that the constraint on appearance probabilities in (1) is easier to understand in the setting of model management than the foregoing constraint on initial acceptance probabilities. The closest solution to the exponential version of R-TBS adapts the weighted sampling algorithm of Chao [7] to batches and forward decay; we call the resulting algorithm B-Chao and describe it in the Online Appendix. Unfortunately, as discussed in the appendix, the relation in (1) is violated both during the initial fill-up phase and whenever the data arrival rate becomes slow relative to the decay rate, so that the sample contains “overweight” items. Including overweight items causes over-representation of older items during initial fill-up or newer items during low inflow, thus potentially degrading predictive accuracy. The root of the issue is that the sample size is nondecreasing over time. The R-TBS algorithm is the first algorithm to correctly (and, for exponential decay, optimally) deal with “underflows” by allowing the sample to shrink—thus handling data streams whose flow rates vary unrestrictedly over continuous time. The current paper also explicitly handles batch arrivals and explores parallel implementation issues. The VarOpt sampling algorithm of Cohen et al. [8]—which was developed to solve the specific problem of estimating “subset sums”—can also be modified via forward decay. The resulting algorithm is more efficient than Chao, but as stated in [8], it has the same statistical properties, and hence does not satisfy (1).

Model management: A key goal of our work is to support model management; see [15] for a survey on methods for detecting changing data—also called “concept drift” in the setting of online learning—and for adapting models to deal with drift. As mentioned previously, one possibility is to re-engineer the learning algorithm. This has been done, for example, with support-vector machines (SVMs) by developing incremental versions of the basic SVM algorithm [6] and by adjusting the training data in an SVM-specific manner, such as by adjusting example weights as in Klittenberg [21]. Klittenberg also considers using curated data selection to learn over concept drift, finding that weighted data selection also improves the performance of learners. Our approach of model retraining using time-biased samples follows this latter approach, and is appealing in that it is simple and applies to a large class of machine-learning models. The recently proposed Velox system for model management [12] ties together online learning and statistical techniques for detecting concept drift. After detecting drift through poor model performance, Velox kicks off batch learning algorithms to retrain the model. Our approach to model management is complementary to the work in [12] and could potentially be used in a system like Velox to help deployed models recover from poor performance more quickly. The developers of the recent MacroBase system [3] have incorporated a time-biased sampling approach to model retraining, for identifying and explaining outliers in fast data streams. MacroBase essentially uses Chao’s algorithm, and so could potentially benefit from the R-TBS algorithm to enforce the inclusion criterion (1) in the presence of highly variable data arrival rates.

8 CONCLUSION

Our experiments with classification and regression algorithms, together with the prior work on graph analytics in [35], indicate the potential usefulness of periodic retraining over time-biased samples to help analytics algorithms deal with evolving data streams without requiring algorithmic re-engineering. To this end we have developed and analyzed several time-biased sampling algorithms that are of independent interest.

In particular, the R-TBS algorithm allows simultaneous control of both the item-inclusion probabilities and the sample size, even when the data arrival rate is unknown and can vary arbitrarily. We have generalized our preliminary algorithms, analyses, and experiments in [18] to arbitrary decay functions. Both theory and empirical results lead us to recommend exponentially and subexponentially decreasing decay functions to achieve reasonable storage and performance. We found that the runtime performance, as well as the resulting accuracy and robustness of ML models, was comparable for the various decay functions that we studied, so we expect that the choice of decay function will be driven by the application setting, as discussed in Section 1. For exponential decay functions, R-TBS maximizes the expected sample size and minimizes sample-size variability. For non-exponential decay functions, we have provided techniques to trade off storage with sample-size behavior in a principled and controllable manner; the user can similarly trade off storage requirements and control of inclusion probabilities.

We have also provided techniques for distributed implementation of R-TBS and T-TBS, and have shown that use of time-biased sampling together with periodic model retraining can improve model robustness in the face of abnormal events and periodic behavior in the data. In settings where (i) the mean data arrival rate is known and (roughly) constant, as with a fixed set of sensors, and (ii) occasional sample overflows can be easily dealt with by allocating extra memory, we recommend use of T-TBS to precisely control item-inclusion probabilities. In many applications, however, we expect that either (i) or (ii) will be violated, in which case we recommend the use of R-TBS. Our experiments showed that R-TBS is superior to sliding windows over a range of parameter values, and hence does not require highly precise parameter tuning; this may be because time-biased sampling avoids the all-or-nothing item inclusion mechanism inherent in sliding windows. In future work, we hope to more fully explore the interplay between model complexity and sampling-algorithm behavior.

Another interesting future direction is to apply and extend our sampling schemes to other types of streaming analytics. An additional goal is to combine our methods with drift-detection techniques to achieve end-to-end model management solutions.

A PROOFS

Proof of Theorem 3.1 Denote by $B_k = |\mathcal{B}_k|$ the (random) size of \mathcal{B}_k for $k \geq 1$. We therefore assume that $\{B_k\}_{k \geq 1}$ are mutually independent and identically distributed as a random variable B having finite mean $b \geq n\gamma$. To prove assertion (i) of the theorem, write

$$C_k = \sum_{i=0}^{k-1} N_{i,k}, \quad (14)$$

where $N_{i,k}$ is the number of sample items from batch \mathcal{B}_{k-i} , i.e., the number of sample items of age $i\Delta$. Observe that the set of sample items of age $i\Delta$ comprises those items in batch \mathcal{B}_{k-i} that survive $k - i + 1$ rounds of Bernoulli sampling with respective success probabilities of $q, p_{i,i+1}, \dots, p_{i,k}$. As is well known, such a sample is probabilistically equivalent to a single Bernoulli sample with success probability $q \times p_{i,i+1} \times \dots \times p_{i,k} = qf_k$. Thus we have

$$\begin{aligned} \mathbb{E}[C_k] &= \mathbb{E}[\mathbb{E}[C_k \mid B_1, \dots, B_k]] = \mathbb{E}\left[\sum_{i=0}^{k-1} \mathbb{E}[N_{i,k} \mid B_{k-i}]\right] \\ &= \mathbb{E}\left[\sum_{i=0}^{k-1} qf_i B_{k-i}\right] = \sum_{i=0}^{k-1} qf_i \mathbb{E}[B_{k-i}] = qbF_{k-1} = nF_{k-1}/F_\infty, \end{aligned}$$

since the $N_{i,k}$ are mutually independent given B_1, \dots, B_k , and each $N_{i,k}$ depends on the batch sizes only through B_{k-i} . Assertion (i) now follows immediately. The proof of assertion (ii) is similar, and uses the fact that $\text{Var}[N_{i,k} \mid B_{k-i}] = qf_i(1 - qf_i)B_{k-i}$.

Assertion (iii) follows from (14) and Hoeffding's inequalities [19]. Indeed, a direct application yields the result in assertion (ii)(a): $\Pr[C_k \geq (1 + \epsilon)n] \leq \exp(-2kn^2\epsilon^2/\bar{b}^2)$ for $\epsilon, k > 0$. To prove assertion (ii)(b), fix $\epsilon > 0$ and $\delta \in (0, \epsilon)$ and observe that, by assertion (i), we have $\mu_k \triangleq \mathbb{E}[C_k] \rightarrow n$ as $k \rightarrow \infty$, so that $\mu_k \geq (1 - \delta)n$ for sufficiently large k . Again applying Hoeffding's inequality, we have that

$$\Pr[C_k \leq (1 - \epsilon)n] \leq \Pr\left[C_k \leq \frac{(1 - \epsilon)}{(1 - \delta)}\mu_k\right] = \Pr[C_k \leq (1 - \epsilon_\delta)\mu_k] \leq e^{-2kn^2\epsilon_\delta^2/\bar{b}^2},$$

where $\epsilon_\delta = (\epsilon - \delta)/(1 - \delta)$.

Proof of Theorem 3.2 Fix f and define the process $\{\mathcal{S}_k\}_{k \geq 0}$ as in Section 3. We claim that $\{\mathcal{S}_k\}_{k \geq 0}$ is an irreducible, aperiodic, time-homogeneous Markov chain with state space $\Sigma = 2^{[0..\bar{b}] \times [0,1,\dots]}$. Indeed the time-homogeneous Markov property follows from the one-step recursive nature of the sample-update process. To prove the rest of the claim, set $\bar{\alpha}(s) = \max\{i : (n, i) \in s\}$ for $s \in \Sigma$, so that $\bar{\alpha}(s)$ is the age of the oldest item(s) in the sample. Next observe that, for any $s, s' \in \Sigma$ there is a positive probability of going from s to \emptyset in one step, and then a positive probability of going from \emptyset to s' in l steps for any $l \geq \bar{\alpha}(s')$.

To prove assertion (i) of the theorem, fix $m \geq 0$. If $\bar{b} = \infty$, so that the batch size is unbounded, then set $s = \{(m', 0)\}$, where $m' = \min\{i \geq m : \Pr[B = i] > 0\}$, and observe that $\Pr[\mathcal{S}_0 = s] \geq q^{m'} \Pr[B = m'] > 0$. If $\bar{b} < \infty$, then set $k = \lceil m/\bar{b} \rceil$ and $s = \{(\bar{b}, 0), \dots, (\bar{b}, k - 1)\}$, and observe that $\Pr[\mathcal{S}_1 = s] \geq \prod_{i=0}^{k-1} (qf_i)^{\bar{b}} \Pr[B = \bar{b}] > 0$.

To prove assertion (ii), it suffices to show that the chain is *recurrent* in that $\Pr[\mathcal{S}_k = s \text{ i.o.}] = 1$ for all $s \in \Sigma$. To this end, we apply an extended version of Foster's Theorem due to Meyn and Tweedie [28, Theorem 2.1(i)]. This result asserts that a sufficient condition for recurrence is the existence of a nonnegative unbounded function V on Σ , a function $k : \Sigma \mapsto \{1, 2, \dots\}$, and a finite subset $A \subseteq \Sigma$ such that

$$\mathbb{E}_s[V(\mathcal{S}_{k(s)})] - V(s) \leq 0 \quad (15)$$

for all $s \in \Sigma \setminus A$, where, in general, $\mathbb{E}_s[g(\mathcal{S}_k)] = \mathbb{E}[g(\mathcal{S}_k) \mid \mathcal{S}_0 = s]$.¹ Let $\psi(s)$ denote the sample size corresponding to state $s \in \Sigma$, i.e., $\psi(s) = \sum_{(n,i) \in s} n$. Then we set $V(s) = (\psi(s) - n)^2 + \bar{\alpha}(s)$. We now develop expressions for the left side of (15), which then determine the required values for $k(s)$ and A . First consider a fixed integer $k \geq 1$ and state $s \in \Sigma$, and observe that we can write $\mathbb{E}_s[V(\mathcal{S}_k)] - V(s) = \mathbb{E}_s[(\psi(\mathcal{S}_k) - n)^2 - (\psi(s) - n)^2] + \mathbb{E}_s[\bar{\alpha}(\mathcal{S}_k) - \bar{\alpha}(s)]$. We analyze each of the two terms on the right separately.

For the first term, denote by n_i the number of age- i sample items when the sample is in state s and set $r_{i,k} = 1 - (f_{i+k}/f_i)$. (Here and elsewhere we suppress the explicit dependence upon f in our notation.) We can write $\psi(\mathcal{S}_k) - \psi(s) = D_1 - D_2$, where $D_1 = \sum_{i=0}^{k-1} N_{i,k}^{(1)}$ and $D_2 = \sum_{i=0}^{\bar{\alpha}(s)} N_{i,k}^{(2)}$, with $N_{i,k}^{(1)}$ and $N_{i,k}^{(2)}$ denoting $\text{Binomial}(B_{k-i}, qf_i)$ and $\text{Binomial}(n_i, r_{i,k})$ random variables, respectively. Here D_1 is the net number of items (after decay) inserted into the sample during the first k steps and D_2 is the total number of initial items that have been removed from the sample during these k steps. Then $D = D_1 - D_2$ is the overall change in the sample size. Observe that $F_\infty^{(2)} \leq F_\infty \leq \sum_{i=0}^\infty if_i < \infty$ by assumption. Recalling that $q = n/(bF_\infty)$, straightforward calculations similar to those given

¹For purposes of analyzing the chain, we let s be any valid state that lies in $\Sigma \setminus A$, even though, when actually sampling, the (random) initial state is $\mathcal{S}_0 = \{(N_0, 0)\}$ where N_0 is a $\text{Binomial}(B, q)$ random variable.

previously show that $E[D^2] = E[D_1^2] - 2E[D_1]E[D_2] + E[D_1^2]$, where

$$\begin{aligned} E[D_1] &= nF_{k-1}/F_\infty, & E[D_1^2] &= nF_{k-1}/F_\infty - n^2F_{k-1}^{(2)}/(bF_\infty^2) + E^2[D_1], \\ E[D_2] &= \sum_{i=1}^{\bar{\alpha}(s)} n_i r_{i,k}, & E[D_2^2] &= \sum_{i=1}^{\bar{\alpha}(s)} n_i r_{i,k}(1 - r_{i,k}) + E^2[D_2]. \end{aligned}$$

Then we have

$$E_s[(\psi(\mathcal{S}_k) - n)^2 - (\psi(s) - n)^2] = E_s[(\psi(s) + D - n)^2 - (\psi(s) - n)^2] = 2E[D](\psi(s) - n) + E[D^2]. \quad (16)$$

Because $\lim_{k \rightarrow \infty} r_{i,k} = 1$ for all i , it follows that

$$\lim_{k \rightarrow \infty} E_s[(\psi(\mathcal{S}_k) - n)^2 - (\psi(s) - n)^2] = -(\psi(s) - n)^2 + \phi_{n,b}, \quad (17)$$

where $\phi_{n,b} = n(1 - nF_\infty^{(2)}/(bF_\infty^2))$.

For the second term, we have $E_s[\bar{\alpha}(\mathcal{S}_k)] = \sum_{l=0}^{k-1} \Pr[\bar{\alpha}(\mathcal{S}_k) > l] = \sum_{l=0}^{k-1} [1 - \Pr[\bar{\alpha}(\mathcal{S}_k) \leq l]]$. The event $\bar{\alpha}(\mathcal{S}_k) \leq l$ occurs if and only if, after k steps, every group—both initial and subsequent—with age $> l$ has lost all of its members due to the decay process. Conditioning on the batch sizes, we have

$$\begin{aligned} \Pr[\bar{\alpha}(\mathcal{S}_k) \leq l] &= E[\Pr[\bar{\alpha}(\mathcal{S}_k) \leq l \mid B_0, \dots, B_{k-1}]] \\ &= E\left[\prod_{i=l+1}^{\bar{\alpha}(s)+k} (1 - qf_i)^{Z_i}\right] \geq \prod_{i=l+1}^{\bar{\alpha}(s)+k} (1 - qf_i)^{\bar{b}} \geq \prod_{i=l+1}^{\bar{\alpha}(s)+k} (1 - f_i)^{\bar{b}} \end{aligned}$$

where $Z_i = B_i$ if $i < k$ and $Z_i = n_{i-k}$ if $i \geq k$. Thus

$$E_s[\bar{\alpha}(\mathcal{S}_k)] \leq \Gamma(\bar{\alpha}(s) + k; \bar{b}), \quad (18)$$

where we define $\Gamma(j; v) = \sum_{l=0}^{j-1} [1 - \prod_{i=l+1}^j (1 - f_i)^v]$ for $j \geq 0$ and $v > 0$. For $v > 0$, denote by v^* the smallest even integer greater than or equal to v . Since $(1 - x)(1 - y) \geq (1 - x - y)$ for any $x, y \in [0, 1]$, we have by induction that

$$\prod_{i=l+1}^j (1 - f_i)^v = \left[\prod_{i=l+1}^j (1 - f_i) \right]^v \geq \left[\prod_{i=l+1}^j (1 - f_i) \right]^{v^*} \geq (1 - F_{l+1,j})^{v^*} \geq 1 - v^* F_{l+1,j},$$

where $F_{a,b} = \sum_{i=a}^b f_i$ and we have used Bernoulli's inequality. It follows that

$$\Gamma(j; v) \leq v^* \sum_{l=0}^{j-1} F_{l+1,j} = v^* \sum_{l=0}^{j-1} \sum_{m=l+1}^{j-1} f_m = v^* \sum_{i=0}^j i f_i \leq v^* \sum_{i=0}^{\infty} i f_i \triangleq \Gamma^*(v) < \infty,$$

for $v > 0$; the second equality follows by interchanging the order of summation. Thus, by (18),

$$E_s[\bar{\alpha}(\mathcal{S}_k) - \bar{\alpha}(s)] \leq \Gamma^*(\bar{b}) - \bar{\alpha}(s) \quad (19)$$

for $k \geq 0$. Combining (17) and (19), we have $\lim_{k \rightarrow \infty} E_s[V(\mathcal{S}_k)] - V(s) \leq -(\psi(s) - n)^2 - \bar{\alpha}(s) + \tilde{\phi}_{n,b}$, where $\tilde{\phi}_{n,b} = \phi_{n,b} + \Gamma^*(\bar{b})$.

Fix $\epsilon > 0$ and, for $s \in \Sigma$, choose $k(s)$ to be large enough so that

$$E_s[V(\mathcal{S}_{k(s)})] - V(s) \leq -(\psi(s) - n)^2 - \bar{\alpha}(s) + \tilde{\phi}_{n,b} + \epsilon.$$

Next choose c and a large enough so that $\min\{(c - n)^2, a\} \geq \tilde{\phi}_{n,b} + \epsilon$, and define the finite set

$$A = \{s \in \Sigma : \psi(s) \leq c \text{ and } \bar{\alpha}(s) \leq a\}. \quad (20)$$

It is now straightforward to verify that (15) holds. Indeed, if $s \in \Sigma \setminus A$, then either $\psi(s) > c$, or $\bar{\alpha}(s) > a$, or both. If $\phi(s) > c$, then, since $-\bar{\alpha}(s) \leq 0$,

$$E_s[V(\mathcal{S}_k)] - V(s) \leq -(\psi(s) - n)^2 + \tilde{\phi}_{n,b} + \epsilon \leq -(c - n)^2 + \tilde{\phi}_{n,b} + \epsilon \leq 0,$$

and if $\bar{\alpha}(s) > a$, then, since $-(\psi(s) - n)^2 \leq 0$,

$$E_s[V(\mathcal{S}_k)] - V(s) \leq -\bar{\alpha}(s) + \tilde{\phi}_{n,b} + \epsilon \leq -a + \tilde{\phi}_{n,b} + \epsilon \leq 0,$$

so that (15) holds in both these cases. If both $\phi(s) > c$ and $\bar{\alpha}(s) > a$, then, of course, (15) holds as well. Assertion (ii) now follows from [28, Theorem 2.1(i)].

To prove assertion (iii), we show that if (8) holds, then the chain is *positive* recurrent in that the expected time between successive visits to any fixed state m is finite. Note that, by the Markov property, the times between successive visits to m are independent and identically distributed. Thus positive recurrence immediately implies assertion (iii)(a). Positive recurrence also implies assertion (iii)(b). Specifically, the chain $\{\mathcal{S}_k\}_{t \geq 0}$ —being irreducible, aperiodic, and positive recurrent—is ergodic, and so has a stationary distribution π [5, Thm. 3.3.1]. This distribution is also a limiting distribution of the chain; in other words, $\mathcal{S}_k \Rightarrow \mathcal{S}_\infty$ for any fixed initial state, where \Rightarrow denotes convergence in distribution and \mathcal{S}_∞ has distribution π [5, Thm. 4.2.1]. Setting $C_\infty = \psi(\mathcal{S}_\infty)$, we then have $C_k \Rightarrow C_\infty$. Moreover, a simple calculation as in the proofs of assertions (i) and (ii) of Theorem 3.1 shows that $E[C_k^2] = qbF_k - q^2bF_k^{(2)} + (qbF_k)^2$, so that $\sup_k E[C_k^2] < \infty$. Thus $\{C_k\}_{k \geq 0}$ is uniformly integrable, and hence $E[C_\infty] = \lim_{k \rightarrow \infty} E[C_k] = n$ [4, p. 338]. Finally, by the strong law of large numbers for Markov chains—see, e.g., [5, Thm. 3.4.1]—we have $\lim_{k \rightarrow \infty} k^{-1} \sum_{i=0}^k C_i = E[C_\infty] = n$ w.p.1.

To show positive recurrence, we establish a stronger drift condition using the function V defined above, namely that

$$E_s[V(\mathcal{S}_{k(s)})] - V(s) \leq -k(s) \quad (21)$$

for $x \in \Sigma \setminus A$, where A has the same structure as in (20), but uses different constants c and a . We then apply Theorem 2.1(ii) in [28]. (The theorem also requires that V be bounded on A , which is obvious.) First observe that our assumption on f implies that, for any $k \geq 1$, we have $\inf_{i \geq 0} r_{i,k} \geq r_k$ for $k \geq 1$, where $r_k = (1 - g_k) \rightarrow 1$ as $k \rightarrow \infty$. Expanding the terms in (16) and using the inequality on $r_{i,k}$, we obtain

$$\begin{aligned} & E_s \left[(\psi(\mathcal{S}_k) - n)^2 - (\psi(s) - n)^2 \right] \\ & \leq 2(\beta_k n - r_k \psi(s)) (\psi(s) - n) + \beta_k n + \beta_k^2 n^2 - 2n\beta_k r_k \psi(s) + \psi(s) + \psi^2(s) \\ & = u_k + v_k \psi(s) - (2r_k - 1)\psi^2(s), \end{aligned} \quad (22)$$

where $\beta_k = F_{k-1}/F_\infty$, $u_k = n\beta_k + n^2\beta_k(\beta_k - 2) \leq 0$, and $v_k = 1 + 2n(\beta_k + r_k - \beta_k r_k)$. Thus

$$E_s[V(\mathcal{S}_k)] - V(s) \leq v_k \psi(s) - (2r_k - 1)\psi^2(s) + \Gamma^*(\bar{b}) - \bar{\alpha}(s)$$

Fix m such that $r_m > 1/2$. From straightforward calculus, we see that

$$v_m \psi(s) - (2r_m - 1)\psi^2(s) + \Gamma^*(\bar{b}) - \bar{\alpha}(s) \leq \frac{v_m}{(8r_m - 4)} + \Gamma^*(\bar{b}) - \bar{\alpha}(s)$$

for all $s \in \Sigma$. Now define A as in (20), choosing c and a so that

$$\max\{v_m c - (2r_m - 1)c^2 + \Gamma^*(\bar{b}), v_m(8r_m - 4)^{-1} + \Gamma^*(\bar{b}) - a\} < -m.$$

The inequality in (21) now follows with $k(s) \equiv m$.

Proof of Theorem 4.1 We first assume that $\pi = \{x^*\}$, so that there exists a partial item in L , and prove the result for $x = x^*$ and then for $x \neq x^*$. We then prove the result when $\pi = \emptyset$.

Proof for $x = x^$:* Observe that when $\pi = \{x^*\}$, we have $\Pr[x^* \in S] = \text{frac}(C)$. First suppose that $\lfloor C' \rfloor = 0$, so that $\text{frac}(C') = C'$. Either the partial item x^* is swapped and ejected in lines 6 and 7 or is retained as a partial item: $\pi' = \{x^*\}$. Thus

$$\begin{aligned} \Pr[x^* \in S'] &= \Pr[x^* \in S' \mid x^* \in L'] \Pr[x^* \in L'] = \text{frac}(C') \Pr[\text{no swap}] = \text{frac}(C')(\text{frac}(C)/C) \\ &= (C'/C) \text{frac}(C) = \theta \Pr[x^* \in S]. \end{aligned}$$

Next suppose that $0 < \lfloor C' \rfloor = \lfloor C \rfloor$. Then the partial item may or may not be converted to a full item via the swap in line 10. Denoting by $r = (1 - (C'/C) \text{frac}(C)) / (1 - \text{frac}(C'))$ the probability that this swap does not occur, we have

$$\begin{aligned} \Pr[x^* \in S'] &= \Pr[x^* \in S' \mid x^* \in \pi'] \Pr[x^* \in \pi'] + \Pr[x^* \in S' \mid x^* \notin \pi'] \Pr[x^* \notin \pi'] \\ &= \text{frac}(C') \cdot \Pr[\text{no swap}] + 1 \cdot \Pr[\text{swap}] = 1 - r(1 - \text{frac}(C')) \\ &= (C'/C) \text{frac}(C) = \theta \Pr[x^* \in S]. \end{aligned}$$

Finally, suppose that $\lfloor C' \rfloor < \lfloor C \rfloor$. Either the partial item x^* is swapped into A in line 14 or ejected in line 17. Thus $\Pr[x^* \in S'] = \Pr[\text{swap}] = (C'/C) \text{frac}(C) = \theta \Pr[x^* \in S]$, establishing the assertion of the lemma for $x = x^*$ when the partial item x^* exists.

Proof for $x \neq x^$:* Still assuming the existence of x^* , set $I_x = 1$ if item x belongs to S' and $Y = I_x = 0$ otherwise. Also set $p_x = \Pr[x \in S'] = \mathbb{E}[I_x]$. Since all full items in S are treated identically, we have $p_x \equiv p$ for $x \in A$, and

$$\mathbb{E}[|S'|] = \mathbb{E}\left[\sum_{x \in A} I_x + I_{x^*}\right] = \sum_{x \in A} \mathbb{E}[I_x] + \mathbb{E}[I_{x^*}] = \lfloor C \rfloor p + p_{x^*}$$

so that, using (9) and the above result,

$$\begin{aligned} \Pr[x \in S'] &= (\mathbb{E}[|S'|] - p_{x^*}) / \lfloor C \rfloor = (C' - (C'/C) \text{frac}(C)) / \lfloor C \rfloor = (C'/C)(C - \text{frac}(C)) / \lfloor C \rfloor \\ &= C'/C = \theta \Pr[x \in S] \end{aligned}$$

for any full item $x \in A$.

Proof when $\pi = \emptyset$: We conclude the proof by observing that, if $\pi = \emptyset$, then $C' = \mathbb{E}[|S'|] = \sum_{x \in A} p_x = \lfloor C \rfloor p = Cp$ and again $\Pr[x \in S'] = C'/C = (C'/C) \Pr[x \in S]$.

Proof of Theorem 4.2 First observe that L is indeed a latent sample: $|A| + |\pi| = \lceil C \rceil$, and $|\pi| \leq 1$. Since $C = C_1 + C_2$ by Line 3, the remainder of assertion (i) of the theorem follows from (9). To prove assertion (ii), observe that for every $x \in A_1$, we have that $x \in A$, so that $\Pr[x \in S] = \Pr[x \in S_1] = 1$. If there is a partial item $x^* \in \pi_1$, we have three cases. If $\text{frac}(C_1) + \text{frac}(C_2) < 1$, then

$$\begin{aligned} \Pr[x^* \in S] &= \Pr[x^* \in \pi] \cdot \Pr[x^* \in S \mid x^* \in \pi] = \frac{\text{frac}(C_1)}{\text{frac}(C_1) + \text{frac}(C_2)} \cdot (\text{frac}(C_1) + \text{frac}(C_2)) \\ &= \text{frac}(C_1) = \Pr[x^* \in S_1]. \end{aligned}$$

If $\text{frac}(C_1) + \text{frac}(C_2) = 1$, then

$$\Pr[x^* \in S] = \Pr[x^* \in A] \cdot \Pr[x^* \in S \mid x^* \in A] = \text{frac}(C_1) \cdot 1 = \Pr[x^* \in S_1].$$

Finally, if $\text{frac}(C_1) + \text{frac}(C_2) > 1$, then

$$\begin{aligned} \Pr[x^* \in S] &= \Pr[x^* \in \pi] \cdot \Pr[x^* \in S \mid x^* \in \pi] + \Pr[x^* \in A] \cdot \Pr[x^* \in S \mid x^* \in A] \\ &= \frac{1 - \text{frac}(C_1)}{2 - \text{frac}(C_1) - \text{frac}(C_2)} \cdot (\text{frac}(C_1) + \text{frac}(C_2) - 1) + \frac{1 - \text{frac}(C_2)}{2 - \text{frac}(C_1) - \text{frac}(C_2)} \cdot 1 \\ &= \text{frac}(C_1) = \Pr[x^* \in S_1]. \end{aligned}$$

This proves assertion (ii), and the proof of assertion (iii) is almost identical.

Proof of Theorem 4.3 Our proof of assertions (i) and (ii) of the theorem is by induction on k . For $k = 1$ and $x \in \mathcal{B}_1$, we see from line 10 that x initially appears in the latent sample $L'_0 = (\mathcal{B}_1, \emptyset, |\mathcal{B}_1|)$ with probability 1. Denote by S'_0 , S_0 , and S_1 random samples generated from L'_0 , L_0 , and L_1 via Algorithm 4 and by C'_0 , C_0 , and C_1 the sample weights of L'_0 , L_0 , and L_1 . By Theorem 4.1, we have $\Pr[x \in S_0] = \rho_1 \Pr[x \in S'_0] = \rho_1 \cdot 1 = \rho_1 f(0) = \rho_1 f(\alpha_{1,1})$. After being unioned with the empty latent sample $(A, \pi, C) = (\emptyset, \emptyset, 0)$, we have by Theorem 4.2 that $\Pr[x \in S_1] = \Pr[x \in S_0] = \rho_1 f(\alpha_{1,1})$. By definition of the downsampling operation, we have $C_0 = \rho_1 |\mathcal{B}_1| = \rho_1 W_1$, and thus, again by Theorem 4.2, $C_1 = C_0 + 0 = \rho_1 W$. This proves (i) and (ii) for $k = 1$. Assume for induction that (i) and (ii) hold for $k - 1$. Then the inductive step for (i) is as given in Section 4.2. To prove the inductive step for (ii), denote by $L_{k,1}$ the latent sample obtained by downsampling L_{k-1} in line 9, by $C_{k,0} = \rho_k |\mathcal{B}_k|$ the sample weight resulting from downsizing $(\mathcal{B}_k, \emptyset, |\mathcal{B}_k|)$ in line 10, and by $C_{k,1}$ the sample weight of $L_{k,1}$. Using Theorem 4.2 and the inductive hypothesis that $\rho_{k-1} = C_{k-1}/W_{k-1}$, we have

$$C_k = C_{k,0} + C_{k,1} = \rho_k |\mathcal{B}_k| + \frac{\rho_k}{\rho_{k-1}} \theta_k C_{k-1} = \rho_k (|\mathcal{B}_k| + \theta_k W_{k-1}) = \rho_k W_k,$$

and the desired result follows.

To prove assertion (iii), fix k and i , and first observe that $f(\alpha_{i,k}) \leq f(\alpha_{i,k-1})$ since f is nonincreasing. If $W_k \geq W_{k-1}$ or $W_{k-1} \leq n$, then $\rho_k \leq \rho_{k-1}$ and the desired result follows. If $W_{k-1} > n$ and $W_k < W_{k-1}$, then we have, setting $\theta_k = e^{-\lambda(t_k - t_{k-1})}$,

$$\begin{aligned} \rho_k f(\alpha_{i,k}) &= \min\left(1, \frac{n}{W_k}\right) f(\alpha_{i,k}) = \min\left(1, \frac{n}{\theta_k W_{k-1} + |\mathcal{B}_k|}\right) \theta_k f(\alpha_{i,k-1}) \leq \frac{n}{W_{k-1}} f(\alpha_{i,k-1}) \\ &= \rho_{k-1} f(\alpha_{i,k-1}). \end{aligned}$$

Proof of Theorem 4.6 We first show that $\tilde{f}_k(\alpha) \leq f(\alpha)$ for all α of the form $\alpha = \alpha_{i,k}$. For $\alpha_{i,k}$ with $i > m(k)$, the assertion follows immediately from the definition of \tilde{f}_k , so we assume that $\alpha = \alpha_{i,k}$ with $i \leq m(k)$ and argue by induction. The assertion is trivially true for $i = m(k)$ by definition of \tilde{f} . Write $m = m(k)$ and assume for induction that the assertion holds for some $i \leq m$. Using the inductive assumption and (12), we have that

$$f(\alpha_{i-1,m}) = f(\alpha_{i,m}) \frac{f(\alpha_{i-1,m})}{f(\alpha_{i,m})} \geq \tilde{f}_k(\alpha_{i,m}) \frac{f(\alpha_{i,m} + \Delta)}{f(\alpha_{i,m})} \geq \tilde{f}_k(\alpha_{i,m}) e^{-\lambda \Delta} = \tilde{f}_k(\alpha_{i-1,m}),$$

and the assertion holds.

To establish assertion (i) of the theorem, observe that $|f(\alpha) - \tilde{f}_k(\alpha)| = f(\alpha) - \tilde{f}_k(\alpha) \leq f(\alpha)$. Because of the check in line 23, a batch of age $\alpha \geq \alpha_k^*$ that joins the consolidated latent sample satisfies $f(\alpha) < \delta_1$. From that point on, the value of α increases for the batch, and hence $f(\alpha)$ decreases, so the relation $f(\alpha) < \delta_1$ continues to hold, proving (i).

We now prove assertion (ii) by induction. For $k = 1$, we have $m(k) = 1$ and thus $\sum_{i=1}^0 |\mathcal{B}_i| f(\alpha_{i,k}) = 0 < \delta_2$. By the induction hypothesis, we have $\sum_{i=1}^{m(k)-1} |\mathcal{B}_i| f(\alpha_{i,k-1}) < \delta_2$. Since f is monotonically decreasing, we have that $\sum_{i=1}^{m(k)-1} |\mathcal{B}_i| f(\alpha_{i,k}) < \delta_2$ prior to line 23 during the processing of the new batch \mathcal{B}_k . It follows that, at line 27, we have $F_\infty - \sum_{i=m(k)}^k f(\alpha_{i,k}) < \delta_2/B_k^*$, where $B_k^* = \max_{i \leq k} |\mathcal{B}_i|$. Since $\sum_{i=1}^{m(k)-1} f(\alpha_{i,k}) + \sum_{i=m(k)}^k f(\alpha_{i,k}) < F_\infty$, we have $\sum_{i=1}^{m(k)-1} f(\alpha_{i,k}) < \delta_2/B_k^*$. Thus $\sum_{i=1}^{m(k)-1} |\mathcal{B}_i| f(\alpha_{i,k}) \leq B_k^* \sum_{i=1}^{m(k)-1} f(\alpha_{i,k}) < \delta_2$.

To prove assertion (iii), recall that $f_i = f(i\Delta)$ and observe that $\sum_{i=m(k)}^k f(\alpha_{i,k}) = \sum_{i=0}^{k-m(k)} f_i$, so that $F_\infty - \sum_{i=m(k)}^k f(\alpha_{i,k}) = \sum_{i=k-m(k)+1}^\infty f_i$. Thus, at line 27, $k - m(k)$ is the smallest integer such

that $\sum_{i=k-m(k)+1}^{\infty} f_i < \delta_2/B_k^*$. Since N is the smallest integer such that $\sum_{i=N}^{\infty} f_i < \delta_2/\bar{b} \leq \delta_2/B_k^*$, we have $k - m(k) + 1 \leq N$, and the assertion follows.

Proof of Proposition 4.7 To prove assertion (i) of the proposition, observe that $n'/W_k^* \geq 1$, and hence $n'/W_i \geq 1$ for all $i \leq k$. Since $\rho_1 = 1$, it follows that $\rho_i = 1$ for $i \leq k$.

To prove assertion (ii), first observe that since f is nonincreasing, we have $\rho_{i,k}^* \geq \rho_{k-1}$ for $i \leq k$ and hence $\rho_k^* \geq \rho_{k-1}$. Now suppose that ρ_k achieves its minimum value in $[t_1, t_k]$ at time $t_{i'}$. Then $\rho_{i'}$ equals 1, $n'/W_{i'}^*$, or $\rho_{i'}^*$. Since $\rho_{i'}^* \geq \rho_{i'-1} \geq \rho_{i'}$, and since $n'/W_j < 1$ for some $t_j \leq t_k$ (so that $\rho_{i'} \leq \rho_j < 1$) we have that $\rho_{i'} = n'/W_{i'} = n'/W_k^*$, the latter equality holding since $\rho_{i'}$ is the smallest value of ρ in $[t_1, t_k]$. Thus for any $i \leq k$ we have $\rho_i \geq \rho_{i'} \geq n'/W_k^*$.

To prove assertion (iii), note that, under the conditions of the proposition, $\rho_k = \rho_k^*$ by (13). As shown above, we have $\rho_k^* > \rho_{k-1}$ —with strict inequality because f is strictly decreasing—and hence $\rho_k > \rho_{k-1}$.

ACKNOWLEDGMENT

The authors wish to thank Valerie Caro for her help with the numerical experiments, as well as the Associate Editor and anonymous reviewers for their very helpful comments.

REFERENCES

- [1] Charu C. Aggarwal. 2006. On biased reservoir sampling in the presence of stream evolution. In *VLDB*. VLDB Endowment, 607–618.
- [2] Noga Alon, Nick Duffield, Carsten Lund, and Mikkel Thorup. 2005. Estimating arbitrary subset sums with few probes. In *PODS*. ACM, 317–325.
- [3] Peter Bailis, Edward Gan, Samuel Madden, Deepak Narayanan, Kexin Rong, and Sahaana Suri. 2017. MacroBase: Prioritizing Attention in Fast Data. In *SIGMOD*. 541–556.
- [4] Patrick Billingsley. 1995. *Probability and Measure* (3rd ed.). Wiley.
- [5] Pierre Brémaud. 1999. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer.
- [6] Gert Cauwenberghs and Tomaso Poggio. 2000. Incremental and Decremental Support Vector Machine Learning. In *NIPS*. 388–394.
- [7] M. T. Chao. 1982. A general purpose unequal probability sampling plan. *Biometrika* (1982), 653–656.
- [8] Edith Cohen, Nick G. Duffield, Haim Kaplan, Carsten Lund, and Mikkel Thorup. 2011. Efficient Stream Sampling for Variance-Optimal Estimation of Subset Sums. *SIAM J. Comput.* 40, 5 (2011), 1402–1431.
- [9] Edith Cohen and Martin J. Strauss. 2006. Maintaining time-decaying stream aggregates. *J. Algo.* 59, 1 (2006), 19–36.
- [10] Graham Cormode, Flip Korn, and Srikanta Tirthapura. 2008. Exponentially decayed aggregates on data streams. In *ICDE*. IEEE, 1379–1381.
- [11] Graham Cormode, Vladislav Shkapenyuk, Divesh Srivastava, and Bojian Xu. 2009. Forward decay: A practical time decay model for streaming systems. In *ICDE*. IEEE, 138–149.
- [12] Daniel Crankshaw, Peter Bailis, Joseph E. Gonzalez, Haoyuan Li, Zhao Zhang, Michael J. Franklin, Ali Ghodsi, and Michael I. Jordan. 2015. The Missing Piece in Complex Analytics: Low Latency, Scalable Model Management and Serving with Velox. In *CIDR*.
- [13] Pavlos S. Efrimidis. 2015. Weighted Random Sampling over Data Streams. In *Algorithms, Probability, Networks, and Games*, Christos D. Zaroliagis, Grammati E. Pantziou, and Spyros C. Koutogiannis (Eds.). Springer, 183–195.
- [14] Pavlos S Efrimidis and Paul G Spirakis. 2006. Weighted random sampling with a reservoir. *Inf. Process. Lett.* 97, 5 (2006), 181–185.
- [15] João Gama, Indre Zliobaite, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *ACM Comput. Surv.* 46, 4 (2014), 44.
- [16] Rainer Gemulla and Wolfgang Lehner. 2008. Sampling time-based sliding windows in bounded space. In *SIGMOD*. 379–392.
- [17] Hiroshi Haramoto, Makoto Matsumoto, Takuji Nishimura, François Panneton, and Pierre L’Ecuyer. 2008. Efficient Jump Ahead for 2-Linear Random Number Generators. *INFORMS Journal on Computing* 20(3) (2008), 385–390.
- [18] Brian Hentschel, Peter J. Haas, and Yuanyuan Tian. 2018. Temporally-Biased Sampling for Online Model Management. In *EDBT*. 109–120.

- [19] Wassily Hoeffding. 1963. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc* 58 (1963), 13–30.
- [20] Ioannis Katakis, Grigorios Tsoumakas, and Ioannis Vlahavas. 2008. An Ensemble of Classifiers for Coping with Recurring Contexts in Data Streams. In *ECAI*. 763–764.
- [21] Ralf Klittenberg. 2004. Learning drifting concepts: Example selection vs. example weighting. *Intell. Data Anal.* 8, 3 (2004), 281–300.
- [22] D. E. Knuth. 1998. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms* (3rd ed.).
- [23] D. B. Lahiri. 1951. A method of sample selection providing unbiased ratio estimates. *Bull. Intl. Statist. Inst.* 33 (1951), 133–140.
- [24] M. Lichman. 2013. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [25] A. I. McLeod and D. R. Bellhouse. 1983. A Convenient Algorithm for Drawing a Simple Random Sample. 32 (1983), 182–184.
- [26] Alexander J. McNeil, Rüdiger Frey, and Paul Embrechts. 2015. *Quantitative Risk Management: Concepts, Techniques and Tools* (second ed.).
- [27] Memcached. 2017. Retrieved 2017-07-13 from <https://memcached.org>
- [28] Sean P. Meyn and R. L. Tweedie. 1994. State-Dependent Criteria for Convergence of Markov Chains. *Ann. Appl. Probab.* 4, 1 (1994), 149–168.
- [29] Frank Olken and Doron Rotem. 1995. Random sampling from databases: a survey. *Statist. Comput.* 5, 1 (1995), 25–42.
- [30] Zhengping Qian, Yong He, Chunzhi Su, Zhuojie Wu, Hongyu Zhu, Taizhi Zhang, Lidong Zhou, Yuan Yu, and Zheng Zhang. 2013. TimeStream: Reliable stream computation in the cloud. In *EuroSys*.
- [31] Redis. 2017. Retrieved 2017-07-13 from <https://redis.io>
- [32] Ernst Stadlober and Heinz Zechner. 1999. The Patchwork Rejection Technique for Sampling from Unimodal Distributions. *ACM Trans. Model. Comput. Simul.* 9, 1 (1999), 59–80.
- [33] Kaggle Sudalai Rajkumar. 2017. An Interactive Deep Dive into the Kaggle Data Science Survey. <https://www.kaggle.com/sudalairajkumar/an-interactive-deep-dive-into-survey-results>
- [34] Andrew Whitmore, Anurag Agarwal, and Li Da Xu. 2015. The Internet of Things – A survey of topics and trends. *Information Systems Frontiers* 17, 2 (2015), 261–274.
- [35] Wenlei Xie, Yuanyuan Tian, Yannis Sismanis, Andrey Balmin, and Peter J. Haas. 2015. Dynamic interaction graphs with probabilistic edge decay. In *ICDE*. 1143–1154.
- [36] Matei Zaharia, Tathagata Das, Haoyuan Li, Timothy Hunter, Scott Shenker, and Ion Stoica. 2013. Discretized Streams: Fault-Tolerant Streaming Computation at Scale. In *SOSP*.