Data Sketches for Disaggregated Subset Sum and Frequent Item Estimation

Daniel Ting
Tableau Software
Seattle, Washington
dting@tableau.com

ABSTRACT

We introduce and study a new data sketch for processing massive datasets. It addresses two common problems: 1) computing a sum given arbitrary filter conditions and 2) identifying the frequent items or heavy hitters in a data set. For the former, the sketch provides unbiased estimates with state of the art accuracy. It handles the challenging scenario when the data is disaggregated so that computing the per unit metric of interest requires an expensive aggregation. For example, the metric of interest may be total clicks per user while the raw data is a click stream with multiple rows per user. Thus the sketch is suitable for use in a wide range of applications including computing historical click through rates for ad prediction, reporting user metrics from event streams, and measuring network traffic for IP flows.

We prove and empirically show the sketch has good properties for both the disaggregated subset sum estimation and frequent item problems. On i.i.d. data, it not only picks out the frequent items but gives strongly consistent estimates for the proportion of each frequent item. The resulting sketch asymptotically draws a probability proportional to size sample that is optimal for estimating sums over the data. For non i.i.d. data, we show that it typically does much better than random sampling for the frequent item problem and never does worse. For subset sum estimation, we show that even for pathological sequences, the variance is close to that of an optimal sampling design. Empirically, despite the disadvantage of operating on disaggregated data, our method matches or bests priority sampling, a state of the art method for pre-aggregated data and performs orders of magnitude better on skewed data compared to uniform sampling. We propose extensions to the sketch that allow it to be used in combining multiple data sets, in distributed systems, and for time decayed aggregation.

CCS CONCEPTS

• **Mathematics of computing** → **Probabilistic algorithms**; • **Theory of computation** → **Sketching and sampling**;

KEYWORDS

Data sketching, subset sum estimation, counting, frequent item, heavy hitters, sampling

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGMOD'18, June 10-15, 2018, Houston, TX, USA
© 2018 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-4703-7/18/06.
<https://doi.org/10.1145/3183713.3183759>

ACM Reference Format:

Daniel Ting. 2018. Data Sketches for Disaggregated Subset Sum and Frequent Item Estimation. In *SIGMOD'18: 2018 International Conference on Management of Data, June 10-15, 2018, Houston, TX, USA*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3183713.3183759>

1 INTRODUCTION

When analyzing massive data sets, even simple operations such as computing a sum or mean are costly and time consuming. These simple operations are frequently performed both by people investigating the data interactively as well as in automated systems which must monitor or collect a multitude of statistics.

Data sketching algorithms enable the information in these massive datasets to be efficiently processed, stored, and queried. This allows them to be applied, for example, in real-time systems, both for ingesting massive data streams or for interactive analyses.

In order to achieve this efficiency, sketches are designed to only answer a specific class of question, and there is typically error in the answer. In other words, it is a form of lossy compression on the original data where one must choose what to lose in the original data. A good sketch makes the most efficient use of the data so that the errors are minimized while having the flexibility to answer a broad range of questions of interest. Some sketches, such as HyperLogLog, are constrained to answer very specific questions with extremely little memory. On the other end of the spectrum, sampling based methods such as coordinated sampling [3], [8] are able to answer almost any question on the original data but at the cost of far more space to achieve the same approximation error.

We introduce a sketch, Unbiased Space Saving, that simultaneously addresses two common data analysis problems: the disaggregated subset sum problem and the frequent item problem. This makes the sketch more flexible than previous sketches that address one problem or the other. Furthermore, it is efficient as it provides state of the art performance on the disaggregated subset sum problem. On i.i.d. streams it has a stronger provable consistency guarantee for frequent item count estimation than previous results, and on non-i.i.d. streams it performs well both theoretically and empirically. In addition, we derive an error estimator with good coverage properties that allows a user to assess the quality of a disaggregated subset sum result.

The disaggregated subset sum estimation is a more challenging variant of the subset sum estimation problem [14], the extremely common problem of computing a sum or mean over a dataset with arbitrary filtering conditions. In the disaggregated subset sum problem [6], [18] the data is "disaggregated" so that a per item metric of interest is split across multiple rows. For example in an ad click stream, the data arrives as a stream of single clicks that are identified with each ad while the metric of interest may be the total

number of clicks per ad. The frequent item problem is the problem of identifying the heavy hitters or most frequent items in a dataset. Several sketches exist for both these individual problems. In particular, the Sample and Hold methods of [6], [16], [18] address the disaggregated subset sum estimation problem. Frequent item sketches include the Space Saving sketch [24], Misra-Gries sketch [25], and Lossy Counting sketch [23].

Our sketch is an extension of the Space Saving frequent item sketch, and as such, has stronger frequent item estimation properties than Sample and Hold. In particular, unlike Sample and Hold, theorem 3 gives both that a frequent item will eventually be included in the sketch with probability 1, and that the proportion of times it appears will be consistently estimated for i.i.d. streams. In contrast to frequent item sketches which are biased, our Unbiased Space Saving sketch gives unbiased estimates for any subset sum, including subsets containing no frequent items.

Our contributions are in three parts: 1) the development of the Unbiased Space Saving sketch, 2) the generalizations obtained from understanding the properties of the sketch and the mechanisms by which it works, and 3) the theoretical and empirical results establishing the correctness and efficiency of the sketch for answering the problems of interest. In particular, the generalizations allow multiple sketches to be merged so that information from multiple data sets may be combined as well as allowing it to be applied in distributed system. Other generalizations include the ability to handle signed and real-valued updates as well as time-decayed aggregation. We empirically test the sketch on both synthetic and real ad prediction data. Surprisingly, we find that it even outperforms priority sampling, a method that requires pre-aggregated data.

This paper is structured as follows. First, we describe the disaggregated subset sum problem, some of its applications, and related sketching problems. We then introduce our sketch, Unbiased Space Saving, as a small but significant modification of the Space Saving sketch. We examine its relation to other frequent item sketches, and show that they differ only in a "reduction" operation. This is used to show that any unbiased reduction operation yields an unbiased sketch for the disaggregated subset sum estimation problem. The theoretical properties of the sketch are then examined. We prove its consistency for the frequent item problem and for drawing a probability proportional to size sample. We derive a variance estimator and show that it can be used to generate good confidence intervals for estimates. Finally, we present experiments using real and synthetic data.

2 TWO SKETCHING PROBLEMS

2.1 Disaggregated subset sum problem

Many data analysis problems consist of a simple aggregation over some filtering and group by conditions.

```
SELECT sum(metric), dimensions
FROM table
WHERE filters
GROUP BY dimensions
```

This problem has several variations that depend on what is known about the possible queries and about the data before the sketch is constructed. For problems in which there is no group by clause and the set of possible filter conditions are known before

the sketch is constructed, counting sketches such as the Count-Min sketch [10] and the AMS, or Count, sketch [2] are appropriate. When the filters and group by dimensions are not known and arbitrary, the problem is the subset sum estimation problem. Sampling methods such as priority sampling [14] can be used to solve it. These work by exploiting a measure of importance for each row and sampling important rows with high probability. For example, when computing a sum, the rows containing large values contribute more to the sum and should be retained in the sample.

The disaggregated subset sum estimation problem is a more difficult variant where there is little to no information about row importance and only a small amount of information about the queries. For example, many user metrics, such as number of clicks, are computed as aggregations over some event stream where each event has the same weight 1 and hence, the same importance. Filters and group by conditions can be arbitrary except for a small restriction that one cannot query at a granularity finer than a specified unit of analysis. In the click example, the finest granularity may be at the user level. One is allowed to query over arbitrary subsets of users but cannot query a subset of a single user's clicks. The data is "disaggregated" since the relevant per unit metric is split across multiple rows. We refer to something at the smallest unit of analysis as an *item* to distinguish it from one row of data.

More formally, the disaggregated subset sum problem answers arbitrary queries of the form

```
SELECT sum(metric), dimensions
FROM (SELECT sum(raw_metric) as metric,
           unit_dimensions
      FROM table
      GROUP BY unit_dimensions
    ) preaggregation
WHERE filters
GROUP BY dimensions
```

given only the unaggregated table and a fixed unit of analysis.

Since pre-aggregating to compute per unit metrics does not reduce the amount of relevant information, it follows that the best accuracy one can achieve is to first pre-aggregate and then apply a sketch for subset sum estimation. This operation, however, is extremely expensive, especially as the number of units is often large. Examples of units include users and ad id pairs for ad click prediction, source and destination IP pairs for IP flow metrics, and distinct search queries or terms. Each of these have trillions or more possible units generating quadrillions or more rows of data.

Several sketches based on sampling have been proposed that address the disaggregated subset sum problem. These include the bottom-k sketch [7] which samples items uniformly at random, the class of "NetFlow" sketches [15], and the Sample and Hold sketches [6], [16], [18]. Of these, the Sample-and-Hold sketches are superior as they use strictly more information than the other methods to construct samples and maintain aggregate statistics. We describe them in more depth in section 4.4.

The Unbiased Space Saving sketch we propose throws away even less information than previous sketches. Despite being computed on disaggregated data, this surprisingly allows it to match or slightly outperform priority sampling, a nearly optimal subset sum estimation algorithm [30], which uses pre-aggregated data.

2.2 Applications

The disaggregated subset sum problem has many applications. These include machine learning and ad prediction [29], analyzing network data [15], [6], detecting distributed denial of service attacks [28], database query optimization and join size estimation [31], as well as analyzing web users' activity logs or other business intelligence applications.

For example, in ad prediction the historical click-through rate and other historical data are among the most powerful features for future ad clicks [19]. Since there is no historical data for newly created ads, one may use historical click or impression data for previous ads with similar attributes such as the same advertiser or product category [27]. In join size estimation, the sketch may be used to estimate the size under the nearly arbitrary filtering conditions that a user might impose.

It also can be naturally applied to hierarchical aggregation problems. For network traffic data, IP addresses are arranged hierarchically. A network administrator may both be interested in individual nodes that receive or generate an excess of traffic or aggregated traffic statistics on a subnet. Several sketches have been developed to exploit hierarchical aggregations including [9], [26], and [32]. Since a disaggregated subset sum sketch can handle arbitrary group by conditions, it can compute the next level in a hierarchy.

2.3 Frequent item problem

The frequent item or heavy hitter problem is related to the disaggregated subset sum problem. Like the disaggregated subset sum problem, frequent item sketches are computed with respect to a unit of analysis which defines a partial aggregation of the data. Rather than allowing for arbitrary subset queries, the problem is more simple as only the most frequent items are of interest. Indeed, if pre-aggregated data were available, the frequent item problem can be solved exactly by using a max priority queue.

Most frequent item sketches are deterministic and have deterministic guarantees on both the identification of frequent items and the error in the counts of individual items. However, since counts in frequent item sketches are biased, further aggregation on the sketch can lead to large errors when bias accumulates, as shown in section 5.3.

Our sketch is based on a deterministic frequent item sketch, but applies randomization to generate unbiased count estimates. This allows it to be used in subset sum queries. Furthermore, theorems 3 and 10 show it retains good frequent item estimation properties.

3 UNBIASED SPACE-SAVING

Our sketch is based on the Space Saving sketch [24] used in frequent item estimation. We refer to it as Deterministic Space Saving to differentiate it from our randomized sketch. For simplicity, we consider the case where the metric of interest is the count for each item. The Deterministic Space Saving sketch works by maintaining a list of m bins labeled by distinct items. A new row with item i increments i 's counter if it is in the sketch. Otherwise, the smallest bin is incremented, and its label is changed to i . Our sketch introduces one small modification. If $\hat{N}_{min}$ is the count for the smallest bin, then only change the label with probability $1/(\hat{N}_{min} + 1)$. This

Algorithm 1 Space-Saving algorithms

- Maintain an m list of $(item, count)$ pairs initialized to have count 0.
- For each new row in the stream, let x_{new} be its item and increment the corresponding counter if the item is in the list. Otherwise, find the pair $(x_{min}, \hat{N}_{min})$ with the smallest count. Increment the counter and replace the item label with x_{new} with probability p .
- For the original Space Saving algorithm $p = 1$. For unbiased count estimates $p = 1/(\hat{N}_{min} + 1)$.

Notation	Definition
t	Number of rows encountered or time
$\hat{N}_i(t)$	Estimate for item i at time t
$\hat{N}_{min}(t)$	Count in the smallest bin at time t
n_i, n_{tot}	True count for item i and total over all items
$\hat{N}_S, n_S$	Estimated and true total count of items in S
$\mathbf{N}, \mathbf{n}$	Vector of estimated and true counts
p_i	Relative frequency n_i/n_{tot} of item i
m	Number of bins in sketch
Z_i	Binary indicator if item i is a label in the sketch
π_i	Probability of inclusion $P(Z_i = 1)$
C_S	Number of items from set S in the sketch

Table 1: Table of symbols

change provably yields unbiased counts as shown in theorem 1. Algorithm 1 describes these Space Saving sketches more formally.

THEOREM 1. *For any item x , randomized Space-Saving given in algorithm 1 gives an unbiased estimate of the count of x .*

PROOF. Let $\hat{N}_x(t)$ denote the estimate for the count of x at time t and $\hat{N}_{min}(t)$ be the count in the smallest bin. We show that the expected increment to $\hat{N}_x(t)$ is 1 if x is the next item and 0 otherwise. Suppose x is the next item. If it is in the list of counters, then it is incremented by exactly 1. Otherwise, it is incremented by $\hat{N}_{min}(t) + 1$ with probability $1/(\hat{N}_{min}(t) + 1)$ for an expected increment of 1. Now suppose x is not the next item. The estimated count $\hat{N}_x(t)$ can only be modified if x is the label for the smallest count. It is incremented with probability $\hat{N}_x(t)/(\hat{N}_x(t) + 1)$. Otherwise $\hat{N}_x(t + 1)$ is updated to 0. This gives the update an expected increment of $\mathbb{E}\hat{N}_x(t + 1) - \hat{N}_x(t) = (\hat{N}_x(t) + 1)\hat{N}_x(t)/(\hat{N}_x(t) + 1) - \hat{N}_x(t) = 0$ when the new item is not x . $\square$

We note that although given any fixed item x , the estimate of its count is unbiased, each stored pair often contains an overestimate of the item's count. This occurs since any item with a positive count will receive a downward biased estimate of 0 conditional on it not being in the sketch. Thus, conditional on an item appearing in the list, the count must be biased upwards.

4 RELATED SKETCHES AND FURTHER GENERALIZATIONS

Although our primary goal is to demonstrate the usefulness of the Unbiased Space-Saving sketch, we also try to understand the

mechanisms by which it works and use this understanding to find extensions and generalizations. Readers only interested in the properties of Unbiased Space Saving may skip to the next section.

In particular, we examine the relationships between Unbiased Space Saving and existing deterministic frequent items sketches as well as its relationship with probability proportional to size sampling. We show that existing frequent item sketches all share the same structure as an exact increment of the count followed by a size reduction. Existing methods implement this size reduction as an adaptive sequential thresholding operation that biases the counts. Our modification replaces the thresholding operation with an unbiased subsampling operation. This also allows us to extend the sketch, such as endowing it with a merge operation that can be used to combine datasets or in distributed computing environments.

The sampling design in the reduction step may also be chosen to give the sketch different properties. For example, time-decayed sampling methods may be used to weight recently occurring items more heavily. If multiple metrics are being tracked, multi-objective sampling [5] may be used.

4.1 Probability proportional to size sampling

Our key observation in generalizing Unbiased Space Saving is that the choice of label is a sampling operation. In particular, this sampling operation chooses the item with probability proportional to its size. We briefly review probability proportional to size sampling and priority sampling as well as the Horvitz-Thompson estimator which unbias the sum estimate from any biased sampling scheme. Probability proportional to size sampling (PPS) is of special importance for sampling for subset sum estimation as it is essentially optimal. Any good sampling procedure mimics PPS sampling.

For unequal probability samples, an unbiased estimator for the sum over the true population $\{x_i\}$ is given by the Horvitz-Thompson estimator $\hat{S} = \sum_i \frac{x_i Z_i}{\pi_i}$ where Z_i denotes whether x_i is in the sample and $\pi_i = P(Z_i = 1)$ is the inclusion probability.

When drawing a sample of fixed size m , it is trivial to see that an optimal set of inclusion probabilities is given by $\pi_i \propto x_i$ when this is possible. In other words, it generates a probability proportional to size (PPS) sample. In this case, each of the m non-zero terms in the sum is constant, and the estimator is exact and has zero variance. When the data is skewed, drawing a truly probability proportional size sample may be impossible for sample sizes greater than 1. For example, given values 1, 1, and 10, any scheme to draw 2 items with probabilities exactly proportional to size has inclusion probabilities bounded by $1/10$, $1/10$, and 1. The expected sample size is at most $12/10 < 2$. In this case, one often chooses inclusion probabilities $\pi_i = \min\{\alpha x_i, 1\}$ for some constant α . The inclusion probabilities are proportional to the size if the size is not too large and 1 otherwise.

Many algorithms exist for generating PPS samples. We briefly describe two as they are necessary for the merge operation given in section 4.5. The splitting procedure of [13] provides a class of methods to generate a fixed size PPS sample with the desired inclusion probabilities. Another method which approximately generates a PPS sample is priority sampling. Instead of exact inclusion probabilities which are typically intractable to compute, priority sampling generates a set of pseudo-inclusion probabilities.

Priority sampling is a method that approximately draws a PPS sample. It generates a random priority $R_i = U_i/n_i$ for an item i with value n_i . The values corresponding to the m smallest priorities form the sample. Surprisingly, by defining the threshold τ be the $(m+1)^{th}$ smallest priority, it can be shown that for almost any function of just the samples, the expected value under this sampling scheme is the same as the expected value under independent $Bernoulli(\min\{1, n_i \tau_i\})$ sampling.

4.2 Misra-Gries and frequent item sketches

The Misra-Gries sketch [25], [12], [21] is a frequent item sketch and is isomorphic to the Deterministic Space Saving sketch [1]. The only difference is that it decrements all counters rather than incrementing the smallest bin when processing an item that is not in the sketch. Thus, the count in the smallest bin for the Deterministic Space Saving sketch is equal to the total number of decrements in the Misra-Gries sketch. Given estimates $\hat{N}$ from a Deterministic Space Saving sketch, the corresponding estimated item counts for the Misra-Gries sketch are $\hat{N}_i^{MG} = (\hat{N}_i - \hat{N}_{min})_+$ where $\hat{N}_{min}$ is the count for the smallest bin and the operation $(x)_+$ truncates negative values to be 0. In other words, the Misra-Gries estimate is the same as the Deterministic Space Saving estimate soft thresholded by $\hat{N}_{min}$. Equivalently, the Deterministic Space Saving estimates are obtained by adding back the total number of decrements $\hat{N}_{min}$ to any nonzero counter in the Misra-Gries sketch.

The sketch has a deterministic error guarantee. When the total number of items is n_{tot} then the error for any item is at most n_{tot}/m .

Other frequent item sketches include the deterministic lossy counting and randomized sticky sampling sketches [23]. We describe only lossy counting as sticky sampling has both worse practical performance and weaker guarantees than other sketches.

A simplified version of Lossy counting applies the same decrement reduction as the Misra-Gries sketch but decrements occur at a fixed schedule rather than one which depends on the data itself. To count items with frequency $> N/m$, all counters are decremented after every m rows. Lossy counting does not provide a guarantee that the number of counters can be bounded by m . In the worst case, the size can grow to $m \log(N/m)$ counters. Similar to the isomorphism between the Misra-Gries and Space-saving sketches, the original Lossy counting algorithm is recovered by adding the number of decrements back to any nonzero counter.

4.3 Reduction operations

Existing deterministic frequent item sketches differ in only the operation to reduce the number of nonzero counters. They all have the form described in algorithm 2 and have reduction operations that can be expressed as a thresholding operation. Although it is isomorphic to the Misra-Gries sketch, Deterministic Space Saving's reduction operation can also be described as collapsing the two smallest bins by adding the larger bin's count to the smaller one's.

Modifying the reduction operation provides the sketch with different properties. We highlight several uses for alternative reduction operations.

The reduction operation for Unbiased Space Saving can be seen as a PPS sample on the two smallest bins. A natural generalization is

Algorithm 2 General frequent item sketching

- Maintain current estimates of counts $\hat{N}(t)$
- Increment $\hat{N}'_{x_{t+1}}(t+1) \leftarrow \hat{N}'_{x_{t+1}}(t) + 1$.
- $\hat{N}(t+1) \leftarrow \text{ReduceBins}(\hat{N}'(t+1), t+1)$

to consider a PPS sample on all the bins. We highlight three benefits of such a scheme. First, items can be added with arbitrary counts or weights. Second, the sketch size can be reduced by multiple bins in one step. Third, there is less quadratic variation added by one sampling step, so error can be reduced. The first two benefits are obvious consequences of the generalization. To see the third, consider when a new row contains an item not in the sketch, and let $\mathcal{J}$ be the set of bins equal to $\hat{N}_{min}$. When using the thresholded PPS inclusion probabilities from section 4.1, the resulting PPS sample has inclusion probability $\alpha = |\mathcal{J}|/(1 + |\mathcal{J}|\hat{N}_{min})$ for the new row's item and $\alpha\hat{N}_{min}$ for bins in $\mathcal{J}$. Other bins have inclusion probability 1. After sampling, the Horvitz-Thompson adjusted counts are $1/|\mathcal{J}| + \hat{N}_{min}$. Unbiased Space Saving is thus a further randomization to convert the real valued increment $1/|\mathcal{J}|$ over $|\mathcal{J}|$ bins to an integer update on a single bin. Since Unbiased Space Saving adds an additional randomization step, the PPS sample has smaller variance. The downside of the more general PPS reduction, however, is that it requires real valued counters that require more space per bin. The update cost when using the stream summary data structure [24] remains $O(1)$.

Changing the sampling procedure can also provide other desirable behaviors. Applying forward decay sampling [11] allows one to obtain estimates that weight recent items more heavily. Other possible operations include adaptively varying the sketch size in order to only remove items with small estimated frequency.

Furthermore, the reduction step does not need to be limited strictly to subsampling. Theorem 2 gives that any unbiased reduction operation yields unbiased estimates. This generalization allows us to analyze Sample-and-Hold sketches.

THEOREM 2. *Any reduction operation where the expected post-reduction estimates are equal to the pre-reduction estimates yields an unbiased sketch for the disaggregated subset estimation problem. More formally, if $\mathbb{E}(\hat{N}(t)|S_{pre}(t)) = \hat{N}_{pre}(t)$ where $S_{pre}(t), \hat{N}_{pre}(t)$ are the sketch and estimated counts before reduction at time step t and $\hat{N}(t)$ is the post reduction estimate, then $\hat{N}(t)$ is an unbiased estimator.*

PROOF. Since $\hat{N}_{pre}(t) = \hat{N}_{post}(t-1) + (n(t) - n(t-1))$, it follows that $\hat{N}(t) - n(t)$ is a martingale with respect to the filtration adapted to $S(t)$. Thus, $\mathbb{E}\hat{N}(t) = n(t)$, and the sketch gives unbiased estimates for the disaggregated subset sum problem. $\square$

We also note that reduction operations can be biased. The merge operation on the Misra-Gries sketch given by [1] performs a soft-thresholding by the size of the $(m+1)^{th}$ counter rather than by 1. This also allows it to reduce the size of the sketch by more than 1 bin at a time. It can be modified to handle deletions and arbitrary numeric aggregations by making the thresholding operation two-sided so that negative values are shrunk toward 0 as well. In this case, we do not provide a theoretical analysis of the properties.

Modifying the reduction operation also yields interesting applications outside of counting. In particular, a reduction operation on matrices can yield accurate low rank decompositions [22], [17].

4.4 Sample and Hold

To our best knowledge, the current state of the art sketches designed to answer disaggregated subset sum estimation problems are the family of sample and hold sketches [18], [16], [6]. These methods can also be described with a randomized reduction operation.

For adaptive sample and hold [6], the sketch maintains an auxiliary variable p which represents the sampling rate. Each point in the stream is assigned a $U_i \sim \text{Uniform}(0, 1)$ random variable, and the items in the sketch are those with $U_i < p$. If an item remains in the sketch starting from time t_0 , then the counter stores the number of times it appears in the stream after the initial time. Every time the sketch becomes too large, the sampling rate is decreased so that under the new rate p' , one item is no longer in the sketch.

It can be shown that unbiased estimates can be obtained by keeping a counter value the same with probability p'/p and decrementing the counter by a random $\text{Geometric}(p')$ random variable otherwise. If a counter becomes negative, then it is set to 0 and dropped. Adding back the mean $(1 - p')/p'$ of the Geometric random variable to the nonzero counters gives an unbiased estimator. Effectively, the sketch replaces the first time an item enters the sketch with the expected $\text{Geometric}(p')$ number of failed tries to enter the sketch. Subsequently, it adds the actual number of times an item appears. Using the memoryless property of Geometric random variables, it is easy to show that the sketch satisfies the conditions of theorem 2. It is also clear that one update step adds more error than Unbiased Space Saving as it potentially adds $\text{Geometric}(p')$ noise with variance $(1 - p')/p'^2$ to every bin. Furthermore, the eliminated bin may not even be the smallest bin. Since p' is the sampling rate, it is expected to be close to 0. By contrast, Unbiased Space Saving has bounded increments of 1 for bins other than the smallest bin, and the only bin that can be removed is the current smallest bin.

The discrepancy is especially prominent for frequent items. A frequent item in an i.i.d. stream for Unbiased Space Saving enters the sketch almost immediately, and the count for the item is nearly exact as shown in theorem 3. For adaptive sample and hold, the first $n_i(1 - p')$ occurrences of item i are expected to be discarded and replaced with a high variance $\text{Geometric}(p')$ random variable. Since p' is typically small in order to keep the number of counters low, most of the information about the count is discarded.

Another sketch, step sample-and-hold, avoids the problem by maintaining counts for each "step" when the sampling rate changes. However, this is more costly both from storage perspective as well as a computational one. For each item in the sketch, computing the expected count takes time quadratic in the number of steps J_i in which the step's counter for the item is nonzero, and storage is linear in J_i .

4.5 Merging and Distributed counting

The more generalized reduction operations allow for merge operations on the sketches. Merge operations and mergeable sketches [1] are important since they allow a collection of sketches, each

Figure 1: In a merge operation, the Misra-Gries sketch simply removes mass from the extra bins with small count. Unbiased Space Saving moves the mass from infrequent items to moderately frequent items. It loses the ability to pick those items as frequent items in order to provide unbiased estimates for the counts in the tail.

answering questions about the subset of data it was constructed on, to be combined to answer a question over all the data. For example, a set of frequent item sketches that give trending news for each country can be combined to give trending news for Europe as well as a multitude of other possible combinations. Another common scenario arises when sketches are aggregated across time. Sketches for clicks may be computed per day, but the final machine learning feature may combine the last 7 days.

Furthermore, merges make sketches more practical to use in real world systems. In particular, they allow for simple distributed computation. In a map-reduce framework, each mapper can quickly compute a sketch, and only a set of small sketches needs to be sent over the network to perform an aggregation at the reducer.

As noted in the previous section, the Misra-Gries sketch has a simple merge operation which preserves its deterministic error guarantee. It simply soft thresholds by the $(m+1)^{th}$ largest counter so that at most m nonzero counters are left. Mathematically, this is expressed as $\hat{N}_i^{new} = (\hat{N}_i^{(1)} + \hat{N}_i^{(2)} - \hat{N}_{(m+1)}^{combined})_+$ where $\hat{N}_i^{(s)}$ is the estimated count from sketch s and $\hat{N}_{(m+1)}^{combined}$ is the $(m+1)^{th}$ smallest nonzero value obtained by summing the estimated counts from the two sketches. Previously, the only merge operations [1], [4] for Deterministic Space Saving were equivalent to first converting it a Misra-Gries sketch to perform a merge, and optionally adding back the threshold value. Theorem 2 shows that by replacing the pairwise randomization with priority sampling or some other sampling procedure still allows one to obtain an Unbiased Space Saving merge that can preserve the expected count in the sketch rather than biasing it downward.

The trade-off required for such an unbiased merge operation is that the sketch may detect fewer of the top items by frequency than the biased Misra-Gries merge. Rather than truncating and preserving more of the "head" of the distribution, it must move mass from the tail closer to the head. This is illustrated in figure 1. We note that this trade-off is not evident in theoretical results.

5 SKETCH PROPERTIES

We study the properties of the space saving sketches here. These include provable asymptotic properties, variance estimates, heuristically and empirically derived properties, behavior on pathological and adversarial sequences, and costs in time and space. In particular, we prove that when the data is i.i.d., the sketch eventually includes all frequent items with probability 1 and that the estimated proportions for these frequent items is consistent. Furthermore, we prove there is a sharp transition between frequent items with consistent estimated proportions and infrequent items which are sampled with probability proportional to their sizes and whose bins contain little to no item specific count information. This is also borne out in the experimental results where the observed inclusion probabilities match the theoretical ones and in estimation error where Unbiased Space Saving matches or even exceeds the accuracy of priority sampling. In pathological cases, we demonstrate that Deterministic Space Saving completely fails at the subset estimation problem. Furthermore, these pathological sequences arise naturally. Any sequence where items' arrival rates change significantly over time forms a pathological sequence. We show that we can derive a variance estimator as well. Since it works under pathological scenarios, the estimator is upward biased. However, we heuristically show that it is close to the variance for a PPS sample. This is confirmed in experiments as well. For both i.i.d. and pathological cases, we examine the resulting empirical inclusion probabilities. Likewise, they behave similarly to a probability proportional to size or priority sample.

5.1 Asymptotic consistency

Our main theoretical result for frequent item estimation states that the sketch contains all frequent items eventually on i.i.d. streams. Thus it does no worse than Deterministic Space Saving asymptotically. We also derive a finite sample bound in section 5.3. Furthermore, the guarantee states that the estimated proportion of times the item appears is strongly consistent and goes to 0. This is better than deterministic guarantees which only ensure that the error is within some constant.

Assume that items are drawn from a possibly infinite, discrete distribution with probabilities $p_1 \geq p_2 \geq \dots$ and, without loss of generality, assume they are labeled by their index into this sequence of probabilities. Let m be the number of bins and t be the number of items processed by the sketch. We will also refer to t as time. Let $\mathcal{I}(t)$ be the set of items that are in the sketch at time t and $Z_i(t) = 1(i \in \mathcal{I}(t))$. To simplify the analysis, we will give a small further randomization by randomly choosing the smallest bin to increment when multiple bins share the same smallest count. Define an absolutely frequent item to be an item drawn with probability $> 1/m$ where m is the number of bins in the sketch. By removing absolutely frequent items and decrementing the sketch size by 1 each time, the set of frequent item can be defined by the condition in corollary 4 which depends only on the tail probability. We first state the theorem and a corollary that immediately follows by induction.

THEOREM 3. *If $p_1 > 1/m$, then as the number of items $t \rightarrow \infty$, $Z_i(t) = 1$ eventually.*

COROLLARY 4. *If $p_i / \sum_{j \geq i} p_j > 1/(m-i+1)$ for all $i < \kappa$ and for some $\kappa < m$, then $Z_i(t) = 1$ for all $i < \kappa$ eventually.*

COROLLARY 5. *Given the conditions of corollary 4, the estimate $\hat{p}_i(t) = \hat{N}_i(t)/t$ is strongly consistent for all $i < \kappa$ as $t \rightarrow \infty$.*

PROOF. Suppose item i becomes sticky after t_0 items are processed. After t_0 , the number of times i appears is counted exactly correctly. As $t \rightarrow \infty$, the number of times i appears after t_0 will dominate the number of times it appears before t_0 . By the strong law of large numbers, the estimate is strongly consistent. $\square$

LEMMA 6. *Let $\alpha = \sum_{j>m} p_j$. For any $\alpha' < \alpha$, $N_{\min}(t) > \alpha' t/m$ eventually as $t \rightarrow \infty$.*

PROOF. Note that any item not in the sketch is added to the smallest bin. The probability of encountering an item not in the sketch is lower bounded by α . Furthermore, by the strong law of large numbers, the actual number of items encountered that are not in the sketch must be $> \alpha' t + m$ eventually. If there are $\alpha' t + m$ items added to the smallest bin, then with m bins, $\hat{N}_{\min}(t) > \alpha' t/m$. $\square$

We now give the proof of theorem 3. The outline of the proof is as follows. We first show that item 1 will always reappear in the sketch if it is replaced. When it reappears, its bin will accumulate increments faster than the average bin, and as long as it is not replaced during this processes, it will escape and never return to being the smallest bin. Since the number of items that can be added before the label on the minimum bin is changed is linear in the size of the minimum bin, there is enough time for item 1 to "escape" from the minimum bin with some constant probability. If it fails to escape on a given try, it will have infinitely many tries, so eventually it will escape.

PROOF. Trivially, $\hat{N}_{\min}(t) \leq t/m$ since there are m bins, and the minimum is less than the average number of items in each bin. If item 1 is not in the sketch, then the smallest bin will take on 1 as its label with probability $p_1/(1 + \hat{N}_{\min}(t)) \geq mp_1/(m + t)$. Since conditional on item 1 not being in the sketch, these are independent events, the second Borel-Cantelli lemma gives that item 1 is in the sketch infinitely often. Whenever item 1 is in the sketch, $\hat{N}_1(t) - t/m$ is a submartingale with bounded increments. Furthermore, it can be lower bounded by an asymmetric random walk $\tilde{N}_1(t) - t/m$ where the expected increment is $\geq p_1 - 1/m$. Let $\epsilon = p_1 - 1/m$. Let t_0 be the time item 1 flips the label of the smallest bin. Lemma 6 gives that the difference $t_0/m - \hat{N}_1(t_0) < t_0(1 - \alpha')/m$ for any $\alpha' < \sum_{j>m} p_j$. If item 1 is not displaced, then after $d = 2t_0(1 - \alpha')/m\epsilon$ additional rows, Azuma's inequality gives after rearrangement, $P(\hat{N}_1(t_0 + d) - (t_0 + d)/m < 0) \leq P(\hat{N}_1(t_0 + d) - \hat{N}_1(t_0) - d/m - d\epsilon < -d\epsilon/2) < \exp(-d\epsilon^2/8) < \exp(-\epsilon(1 - \alpha')/4m)$. The probability that item 1 is instead displaced during this time is $< d/(d + \alpha' t_0)$ which can be simplified to some positive constant that does not depend on t_0 . In other words, there is some constant probability γ such that item 1 will go from being in the smallest bin to a value greater than the mean. From there, there is a constant probability that the bounded sub-martingale $\hat{N}_i(t_0 + d + \Delta) - (t_0 + d + \Delta)/m$ never crosses back to zero or below. Since item 1 appears infinitely often, it must either become sticky or there are infinitely many 0 upcrossing for $\hat{N}_1(t) - t/m$. In the latter case, there is a constant probability $\rho > 0$ that lower bounds the probability the item becomes sticky. Thus a geometric random variable lower bounds the number of tries before item i "sticks," and it must eventually be sticky. $\square$

5.2 Approximate PPS Sample

We prove that for i.i.d. streams, Unbiased Space Saving approximates a PPS sample and does so without the expensive pre-aggregation step. This is born out by simulations as, surprisingly, it often empirically outperforms priority sampling from computationally expensive, pre-aggregated data. Since frequent items are included with probability 1, we consider only the remaining bins and the items in the tail.

LEMMA 7. *Let B_i denote the count in the i^{th} bin. If $p_1 < 1/m$ then $B_i(t) - t/m < (\log t)^2 + 1$ eventually.*

PROOF. If $B_i(t) > t/m$ then $B_i(t)$ is not the smallest bin. In this case, the expected difference after 1 time step is bounded above by $\delta := p_1 - 1/m < 0$. Consider a random walk $W(t)$ with an increment of $1 - 1/m$ with probability p_1 and $-1/m$ otherwise. By Azuma's inequality, if it is started at time $t - s$ at value 1 then the probability it exceeds $(\log t)^2$ is bounded by $P(W(t) - s/m - 1 > c(t) + \delta s) < \exp(-(c(t) + \delta s)^2/2s)$. Since for $B_i(t) - t/m$ to be $> c(t)$, it must upcross 0 at some time $t - s$, maximizing over s gives an upper bound on the probability $B_i(t) - t/m > c(t)$. It is easy to derive that $s = c(t)/\delta$ is the maximizer and the probability is bounded by $\exp(-\delta c(t))$. When $c(t) = (\log t)^2$, $\sum_{t=1}^{\infty} \exp(-\delta c(t)) < \infty$, and the conclusion holds by the Borel-Cantelli lemma. $\square$

LEMMA 8. *If $p_1 < 1/m$ then $0 \leq t/m - \hat{N}_{\min} \leq m(\log t)^2 + m$ and $0 \leq \hat{N}_{\max} - t/m \leq (\log t)^2 + 1$ eventually*

PROOF. Since there are finitely many bins, by the lemma 7, $0 \leq \hat{N}_{\max} - t/m \leq (\log t)^2 + 1$ eventually. The other inequality holds since $t/m - \hat{N}_{\min} < m(\hat{N}_{\max} - t/m)$ $\square$

THEOREM 9. *If $p_1 < 1/m$, then the items in the sketch converge in distribution to a PPS sample.*

PROOF. The label in each bin is obtained by reservoir sampling. Thus it is a uniform sample on the rows that go into that bin. Since all bins have almost exactly the same size $t/m + O((\log t)^2)$, it follows that item i is a label with probability $p_i + O((\log t)^2/t)$. $\square$

The asymptotic analysis of Unbiased Space Saving splits items into two regimes. Frequent items are in the sketch with probability 1 and the associated counts are nearly exact. The threshold at which frequent and infrequent items are divided is given by corollary 4 and is the same as the threshold in the merge operation shown in figure 1. The counts for infrequent items in the tail are all $\hat{N}_{\min}(t)(1 + o(1))$. The actual probability for the item in the bin is irrelevant since items not in the sketch will force the bin's rate to catch up to the rate for other bins in the tail. Since an item changes the label of a bin with probability $1/B$ where B is the size of the bin, the bin label is a reservoir sample of size 1 for the items added to the bin. Thus, the labels for bins in the tail are approximately proportional to their frequency. Figure 2 illustrates that the empirical inclusion probabilities match the theoretical ones for a PPS sample. The item counts are chosen to approximate a rounded *Weibull*(5 $\times$ 10⁵, 0.15) distribution. This is a skewed distribution where the standard deviation is roughly 30 times the mean.

Figure 2: The inclusion probability for each item empirically behaves like the inclusion probability for a probability proportional to size sample. This is also nearly equivalent to the inclusion probabilities for a priority sample.

We note, however, that the resulting PPS sample has limitations not present in PPS samples on pre-aggregated data. For pre-aggregated data, one has both the original value x_i and the Horvitz-Thompson adjusted value x_i/π_i where π_i is the inclusion probability. This allows the sample to compute non-linear statistics such as the variance which is a quadratic function $\sum_i x_i^2 Z_i/\pi_i$. With the PPS samples from disaggregated subset sum sketching, only the adjusted values are observed.

5.3 Pathological sequences

Deterministic Space Saving has remarkably low error when estimating the counts of frequent items [9]. However, we will show that it fails badly when estimating subset sums when the data stream is not i.i.d.. Unbiased Space Saving performs well on both i.i.d. and on pathological sequences.

Pathological cases arise when an item's arrival rate changes over time rather than staying constant. Consider a sketch with 2 bins. For a sequence of c 1's and c 2's followed by a single 3 and 4, the Deterministic Space Saving algorithm will always return 3 and 4, each with count $c + 1$. By contrast, Unbiased Space Saving will return 1 and 2 with probability $(1 - 1/c)^2 \approx 1$ when c is large. Note that in this case, the count for each frequent item is slightly below the threshold that guarantees inclusion in the sketch, $c < n/2$. This example illustrates the behavior for the deterministic algorithm. When an item is not in the "frequent item head" of the distribution then the bins that represent the tail pick the labels of the most recent items without regard to the frequency of older items.

We note that natural pathological sequences can easily occur. For instance, partially sorted data can naturally lead to such pathological sequences. This can occur from sketching the output of some join. Data partitioned by some key where the partitions are processed in order is another case. We explore this case empirically in section 6. Periodic bursts of an item followed by periods in which its frequency drops below the threshold of guaranteed inclusion are another example. The most obvious pathological sequence is the case where every row is unique. The Deterministic Space Saving sketch always consists of the last m items rather than a random sample, and no meaningful subset sum estimate can be derived.

For Unbiased Space Saving, we show that even for non-i.i.d. sequences, it essentially never has an inclusion probability worse than simple random sampling which has inclusion probability $1 - (n_{tot} - n_i)_m / (n_{tot})_m \approx 1 - (1 - n_i/n_{tot})^m$ where $(x)_m$ denotes the m^{th} falling factorial.

THEOREM 10. *An item i occurring n_i times has worst case inclusion probability $\pi_i \geq 1 - (1 - n_i/n_{tot})^m$. An item with asymptotic frequency $n_i = \alpha n/m + o(n/m)$ has an inclusion probability $\pi_i \geq 1 - e^{-\alpha} + o(1)$ as $n, m \rightarrow \infty$.*

PROOF. Whether an item is in the sketch depends only on the sequence of additions to the minimum sized bin. Let T_b be last time an item is added to bin b while it is the minimum bin. Let $C_{i,b}$ be the number of times item i is added to bin b by time T_b and L_b be the count of bin b at time T_b . Item i is not the label of bin b with probability $1 - C_{i,b}/L_b$, and it is not in the sketch with probability $\prod_b (1 - C_{i,b}/L_b)$. Note that for item i to not be in the sketch, the last occurrence of i must have been added to the minimum sized bin. Thus, maximizing this probability under the constraints that $\sum_b L_b \leq n$ and $\sum_b C_{i,b} = n_i$ gives an upper bound on $1 - \pi_i$ and yields the stated result. $\square$

We note that the bound is often quite loose. It assumes a pathological sequence where the minimum sized bin is as large as possible, namely $L_b = n_{tot}/m$. If $L_b \leq \gamma n_{tot}/m$, the asymptotic bound would be $\pi_i \geq 1 - e^{-\alpha/\gamma} + o(1)$.

At the same time, we note that the bound is tight in the sense that one can construct a pathological sequence that achieves the upper bound. Consider the sequence consisting of $n_{tot} - n_i$ distinct items followed by item i for n_i times with n_i and n_{tot} both being multiples of m . It is easy to see that the only way that item i is not in the sketch is for it no bin to ever take on label i and for the bins to all be equal in size to the minimum sized bin. The probability of this event is equal to the given upper bound.

Although Deterministic Space Saving is poor on pathological sequences, we note that if data arrives in uniformly random order or if the data stream consists of i.i.d. data, one expects the Deterministic Space Saving algorithm to share similar unbiasedness properties as the randomized version as in both cases the label for a bin can be treated roughly as a uniform random choice out of the items in that bin.

5.4 Variance

In addition to the count estimate, one may also want an estimate of the variance. In the case of i.i.d. streams, this is simple since it forms an approximate PPS sample. Since the inclusion of items is negatively correlated, a fixed size PPS sample of size m has variance upper bounded by

$$\text{Var}_{PPS}(\hat{N}_i) \leq \alpha n_i (1 - \pi_i). \quad (1)$$

When the marginal sampling probabilities $\pi_i = \min\{1, \alpha n_i\}$ are small, this upper bound is nearly exact. For the non-i.i.d. case, we provide a coarse upper bound. Since $\hat{N}_i(t) - n_i(t)$ is a martingale as shown in theorem 2, the quadratic variation process taking the squares of the increments $\sum_t (\hat{N}_i(t+1) - \hat{N}_i(t) - n_i(t+1) + n_i(t))^2$ yields an unbiased estimate of the variance. There are only two cases where the martingale increment is non-zero: the new item

is i and i is not in the sketch or the new item is not i and item i is in the smallest bin. In each case the expected squared increment is $\hat{N}_{min}(t) - 1$ since the updated value is $1 + \hat{N}_{min}(t)\tilde{Z}_t$ where $\tilde{Z}_t \sim \text{Bernoulli}(N_{min}(t)^{-1})$. Let τ_i be the time when item i becomes "sticky." That is the time at which a bin acquires label i and never changes afterwards. If item i does not become sticky, then $\tau_i = n$. Define $\kappa_i = n_i(\tau_i)$. It is the number of times item i is added until it becomes sticky. This leads to the following upper bound on the variance

$$\text{Var}(\hat{N}_i) \leq \sum_{j=0}^{\kappa_i} \mathbb{E} \left[\left(\hat{N}_{min} - \left\lfloor \frac{j}{m} \right\rfloor \right)_+ - 1 \right] \quad (2)$$

$$\leq \mathbb{E}(\hat{N}_{min}\kappa_i). \quad (3)$$

We note that the same variance argument holds when computing a further aggregation to estimate $n_S = \sum_{i \in S} n_i$ for a set of items S . In this case κ_S is the total number of times items in S are added to the sketch excluding the deterministic additions to the final set of "sticky" bins.

To obtain a variance estimate for a count, we plug in an estimate for $\hat{\kappa}_i$ into equation 3. We use the following estimate

$$\hat{\kappa}_S = \hat{N}_{min} C_S \quad (4)$$

$$\widehat{\text{Var}}(\hat{N}_S) = \hat{N}_{min}^2 C_S \quad (5)$$

where C_S is the greater of 1 and the number of bins labeled with an item in S .

The estimate $\hat{\kappa}_S$ is an upward biased estimate for κ_S . For items with count $\leq \hat{N}_{min}$, one has no information about their relative frequency compared to other infrequent items. Thus, we choose the worst case as our estimate $\hat{\kappa}_S = \hat{N}_{min}$. For items with count $> \hat{N}_{min}$, we also take a worst case approach for estimating κ . Consider a bin with size $\leq V - 1$. The probability that an additional item will cause a change in the label is $1/V$. Since $\hat{N}_{min}$ is the largest possible "non-sticky" bin, it follows $\kappa_i - 1 < Y$ where $Y \sim \text{Geometric}(1/\hat{N}_{min})$. Taking the expectation given $\hat{N}_{min}$ gives the upward biased estimate $\hat{\kappa}_i = \hat{N}_{min} + 1$. In this case, we drop the 1 for simplicity and because it is an overestimate.

We compare this variance estimate with the variance of a Poisson PPS sample and show that they are similar. They are nearly identical for infrequent items, but Unbiased Space Saving adds an additional term to the variance for frequent items. In the i.i.d. scenario for Unbiased Space-Saving, $\mathbb{E}C_i = \pi_i \rightarrow n_i/\alpha$ and $\hat{N}_{min}$ converges to α for some $\alpha > 0$. Plugging these into equation 5 gives a variance estimate of αn_i which differs only by a factor of $1 - \pi_i$ from the variance of a Poisson PPS sample given in equation 1. For infrequent items, π_i is typically small. For frequent items, a Poisson PPS sample has inclusion probability 1 and zero variance. In this case, the worst case behavior for Unbiased Space Saving contributes the same variance as an infrequent item. The similar behavior to PPS samples is also borne out by experimental results. Figure 9 shows that the variance estimate is often very accurate and close to the variance of a true PPS sample.

While the empirical variance estimate in equation 5 provides more accurate results, note that $\hat{N}_{min} \leq n_{tot}/m$ and $\kappa_S \leq n_S$, so that $\text{Var}(\hat{N}_S) \leq n_{tot} n_S / m$ provides a trivial but loose upper bound on the variance that scales inversely with the number of samples m as expected.

5.5 Confidence Intervals

As the inclusion of a specific item is a binary outcome, confidence intervals for individual counts are meaningless. However, the variance estimate allows one to compute Normal confidence intervals when computing sufficiently large subset sums. Thus, a system employing the sketch can provide estimates for the error along with the count estimate itself. These estimates are valid even when the input stream is a worst case non-i.i.d. stream. Experiments in section 6 shows that these Normal confidence intervals have close to or better than advertised coverage whenever the central limit theorem applies, even for pathological streams.

5.6 Robustness

For the same reasons it has much better behavior under pathological sequences, Unbiased Space Saving is also more robust to adversarial sequences than Deterministic Space Saving. Theorem 11 shows that by inserting an additional n_{tot} items, one can force all estimated counts to 0, including estimates for frequent items, as long as they are not too frequent. This complete loss of useful information is a strong contrast to the theoretical and empirical results for Unbiased Space Saving which suggest that polluting a dataset with n_{tot} noise items will simply halve the sample size, since it will still return a sample that approximates a PPS sample.

THEOREM 11. *Let $\mathbf{n}$ be a vector of v counts with $n_{tot} = \sum_{i=1}^v n_i$ and $n_i < 2n_{tot}/m$ for all $i \leq v$. There is a sequence of $2n_{tot}$ rows such that item i appears exactly n_i times, but the Deterministic Space Saving sketch returns an estimate of 0 for all items $i \leq v$.*

PROOF. Sort the items from most frequent to least frequent. This ensures no frequent item will have an estimated count greater than its true count. Add n_{tot} additional distinct items. The resulting deterministic sketch will consist only of the additional distinct items and each bin will have count $2n_{tot}/m \pm 1$. $\square$

5.7 Running time and space complexity

The update operation is identical to the Deterministic Space Saving update except that it changes the label of a bin less frequently. Thus, each update can be performed in $O(1)$ time when the stream summary data structure [24] is used. In this case the space usage is $O(m)$ where m is the number of bins.

6 EXPERIMENTS

We perform experiments with both simulations and real ad prediction data. For synthetic data, we consider three cases: randomly permuted sequences, realistic pathological sequences for Deterministic Space Saving, and "pathological" sequences for Unbiased Space Saving. For each we draw the count for each item using a Weibull distribution that is discretized to integer values. That is $n_i \sim \text{Round}(\text{Weibull}(k, \alpha))$ for item i . The discretized Weibull distribution is a generalization of the geometric distribution that allows us to adjust the tail of the distribution to be more heavy tailed. We choose it over the Zipfian or other truly heavy tailed distributions as few real data distributions have infinite variance. Furthermore, we expect our methods to perform better under heavy tailed distributions with greater data skew as shown in figure 3.

Figure 3: The sketch accuracy improves when the skew is higher and when more and larger bins are contained in the subset. The number of bins is 200.

For more easily reproducible behavior we applied the inverse cdf method $n_i = F^{-1}(U_i)$ where the U_i are on a regular grid of 1000 values rather than independent $Uniform(0, 1)$ random variables. Randomly permuting the order in which individual rows arrive yields an exchangeable sequence which we note is identical to an i.i.d. sequence in the limit by de Finetti's theorem. In each case, we draw at least 10,000 samples to estimate the error.

For real data, we use a Criteo ad click prediction dataset¹. This dataset provides a sample of 45 million ad impressions. Each sample includes the outcome of whether or not the ad was clicked as well as multiple integer valued and categorical features. We did not randomize the order in which data arrives. We pick a subset of 9 of these features. There are over 500 million possible tuples on these features and many more possible filtering conditions.

The Criteo dataset provides a natural application of the disaggregated subset sum problem. Historical clicks are a powerful feature in click prediction [27], [20]. While the smallest unit of analysis is the *ad* or the (*user*, *ad*) pair, the data is in a disaggregated form with one row per impression. Furthermore, since there may not be enough data for a particular *ad*, the relevant click prediction feature may be the historical click through rate for the advertiser or some other higher level aggregation. Past work using sketches to estimate these historical counts [29] include the CountMin counting sketch as well as the Lossy Counting frequent item sketch.

To simulate a variety of possible filtering conditions, we draw random subsets of 100 items to evaluate the randomly permuted case. As expected, subsets which mostly pick items in the tail of the distribution and have small counts also have estimates with higher relative root mean squared error. The relative root mean squared error (RRMSE) is defined as $\sqrt{MSE}/n_S$ where n_S is the true subset sum. For unbiased estimators this is equivalent to σ_S/n_S where σ_S is the standard deviation of the estimator. Note that an algorithm with c times the root mean squared error of a baseline algorithm typically requires c^2 times the space as the variance, not the standard deviation, scales linearly with size.

We compare our method to uniform sampling of items using the bottom-k sketch, priority sampling, and Deterministic Space Saving. Although we do not directly compare with sample and hold methods, we note that figure 2 in [6] shows that sample and hold performs significantly worse than priority sampling.

¹<http://labs.criteo.com/2014/02/kaggle-display-advertising-challenge-dataset/>

Figure 4: Unbiased Space Saving performs orders of magnitude better than uniform sampling of items (Bottom-k) in the smoothed plot of relative error versus the true count. With 100 bins, the error is higher than with 200 bins given in figure 3 but the curve is qualitatively similar.

Figure 5: Unbiased Space Saving performs slightly better than priority sampling on the synthetic data despite priority sampling using pre-aggregated data rather than the raw unaggregated data stream.

Surprisingly, figure 5 shows Unbiased Space Saving performs slightly better than priority sampling even though priority sampling is applied on pre-aggregated data. We are unsure as to the exact reason for this. However, note that, unlike Unbiased Space Saving, priority sampling does not ensure the total count is exactly correct. A priority sample of size 100 when all items have the same count will have relative error of $\approx 10\%$ when estimating the total count.

This added variability in the threshold and the relatively small sketch sizes for the simulations on i.i.d. streams may explain why Unbiased Space Saving performs even better than what could be considered close to a "gold standard" on pre-aggregated data.

6.1 Pathological cases and variance

For pathological sequences we find that Unbiased Space Saving performs well in all cases while Deterministic Space Saving has unacceptably large errors even for reasonable non-i.i.d. sequences. First we consider a pathological sequence for Deterministic Space Saving. This sequence consists of two halves. Each half is an independent i.i.d. stream from a discretized Weibull frequency distribution. This is a natural scenario as the data may be randomly partitioned into blocks, for example, by hashed user id, and each block is fed into the sketch for summarization. As shown in figure 7, Deterministic Space Saving completely ignores infrequent items in the first half of the stream, resulting in large bias and error. In this

Figure 6: The Unbiased Space Saving sketch is able to accurately compute 1 and 2 way marginals on the Criteo dataset. The average relative mse for a marginal count that is between 100k and 200k is $< 5\%$ and for marginals containing more than half the data, the mean squared error drops to under 0.5%. It performs similarly to priority sampling.

Figure 7: Left: Items 1 to 1,000 only appear in the first half of the stream. The inclusion probabilities for a pathological sequence still behave like a PPS sample for Unbiased Space Saving, but only the frequent items in the first half are sampled under Deterministic Space Saving. Right: As a result, Deterministic Space Saving is highly inaccurate when querying items in the first half of the stream.

case, the simulations use small sketches with only 100 bins. The disparity would increase with larger sketches and streams where the bias of Deterministic Space Saving remains the same but the variance decreases for Unbiased Space Saving.

The types of streams that induce worst case behavior for Deterministic and Unbiased Space Saving are different. For Unbiased Space Saving, we consider a sorted stream arranged in ascending order by frequency. Note that the reverse order where the largest items occur first gives an optimally favorable stream for Unbiased Space Saving. Every frequent item is deterministically included in the sketch, and the count is exact. The sequence consists of 10^5 distinct items and 10^9 rows where the item counts are from a discretized Weibull distribution. We use 10,000 bins in these experiments. To evaluate our method, we partition the sequence into 10 epochs containing an equal number of distinct items and estimate the counts of items from each epoch. We find in this case our variance estimate given in equation 5 yields an upward biased estimate of the variance as expected. Furthermore, it is accurate except for very small counts and the last items in a stream. Figure 8 shows the

Figure 8: Left: For a pathological sorted sequence, the true counts are given with bars indicating the average 95% confidence interval width. For epochs 1 to 4, the intervals are truncated below as they extend past 0. Right: Normal confidence intervals generally deliver higher than advertised coverage. The exceptions lie in a regime where the variance estimate is accurate as shown in figure 9, but the sample contains too few items from the epoch to apply the central limit theorem.

true counts and the corresponding 95% confidence intervals computed as $\hat{N}_S \pm 1.96\widehat{\text{Var}}(\hat{N}_S)$. In epochs 4 and 5, there are on average roughly 3 and 13 items in the sample, and the asymptotic properties from the central limit theorem needed for accurate normal confidence intervals have not or are not fully manifested. For epochs 1 and 2, the upward bias of the variance estimate gives 100% coverage despite the central limit theorem not being applicable. The coverage of a confidence interval is defined to be the probability the interval includes the true value. A 95% confidence interval should have almost exactly 95% coverage. Lower coverage represents an underestimation of variability or risk. Less harmful is higher coverage, which represents an overly conservative estimation of variability.

We note that the behavior of Deterministic Space Saving is easy to derive in this case. The first 9 epochs have estimated count of 0 and the last epoch has estimated count $n_{tot} = 10^9$. Figure 10 shows that except for small counts, Unbiased Space Saving performs an order of magnitude better than Deterministic Space Saving.

7 CONCLUSION

We have introduced a novel sketch, Unbiased Space Saving, that answers both the disaggregated subset sum and frequent item problems and gives state of the art performance under all scenarios. Surprisingly, for the disaggregated subset sum problem, the sketch can outperform even methods that run on pre-aggregated data.

We prove that asymptotically, it can answer the frequent item problem for i.i.d. sequences with probability 1 eventually. Furthermore, it gives stronger probabilistic consistency guarantees on the accuracy of the count than previous results for Deterministic Space Saving. For non-i.i.d. streams, we show that Unbiased Space Saving still has attractive frequent item estimation properties and exponential concentration of inclusion probabilities to 1.

For the disaggregated subset sum problem, we prove that the sketch provides unbiased results. For i.i.d. stream, we show that items selected for the sketch are sampled approximately according

Figure 9: Left: For pathological streams, the estimated standard deviation from equation 5 is shown to be accurate and match the true standard deviation for counts that are not too large or small. Right: Even for pathological streams, the variance closely matches the variance from a PPS sample.

Figure 10: Deterministic Space Saving gives grossly incorrect answers on all epochs. For epochs 9 and 10, the error is 50x that of Unbiased Space Saving. For extremely small counts constituting $< 0.002\%$ of the total, the possibility of overestimation causes Unbiased Space Saving to have worse error compared to Deterministic Space Saving which always estimates 0.

to an optimal PPS sample. For non-i.i.d. streams we show that it empirically performs well and is close to a PPS sample even if given a pathological stream for which Deterministic Space Saving fails badly on. We derive a variance estimator for subset sum estimation and show that it is nearly equivalent to the estimator for a PPS sample. It is shown to be accurate on pathological sequences and yields confidence intervals with good coverage.

We study Unbiased Space Saving's behavior and connections to other data sketches. In particular, we identify the primary difference between many of the frequent item sketches is a slightly different operation to reduce the number of bins. We use that understanding to provide multiple generalizations to the sketch which allow it to be applied in distributed settings, handle weight decay over time, and adaptively change its size over time. This also allows us to compare Unbiased Space Saving to the family of sample and hold sketches that are also designed to answer the disaggregated subset

sum problem. This allows us to also mathematically show that Unbiased Space Saving is superior.

REFERENCES

- [1] P. K. Agarwal, G. Cormode, Z. Huang, Jeff M Phillips, Z. Wei, and K. Yi. 2013. Mergeable summaries. *ACM Transactions on Database Systems* 38, 4 (2013), 26.
- [2] N. Alon, Y. Matias, and M. Szegedy. 1999. The Space Complexity of Approximating the Frequency Moments. *J. Comput. System Sci.* 58, 1 (1999), 137–147.
- [3] K. R.W. Brewer, L.J. Early, and S.F. Joyce. 1972. Selecting several samples from a single population. *Australian & New Zealand Journal of Statistics* 14, 3 (1972), 231–239.
- [4] M. Cafaro, M. Pulimeno, and P. Tempesta. 2016. A parallel space saving algorithm for frequent items and the hurwitz zeta distribution. *Information Sciences* 329 (2016), 1–19.
- [5] E. Cohen. 2015. Multi-objective weighted sampling. In *Hot Topics in Web Systems and Technologies (HotWeb), 2015 Third IEEE Workshop on*. IEEE, 13–18.
- [6] E. Cohen, N. Duffield, H. Kaplan, C. Lund, and M. Thorup. 2007. Sketching unaggregated data streams for subpopulation-size queries. In *PODS*.
- [7] Edith Cohen and Haim Kaplan. 2007. Summarizing data using bottom-k sketches. In *PODC*.
- [8] E. Cohen and H. Kaplan. 2013. What You Can Do with Coordinated Samples. In *RANDOM*.
- [9] G. Cormode and M. Hadjieleftheriou. 2008. Finding frequent items in data streams. *VLDB* (2008).
- [10] G. Cormode and S. Muthukrishnan. 2005. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms* 55, 1 (2005), 58–75.
- [11] G. Cormode, V. Shkapenyuk, D. Srivastava, and B. Xu. 2009. Forward decay: A practical time decay model for streaming systems. In *ICDE*. IEEE, 138–149.
- [12] E. D. Demaine, A. López-Ortiz, and J. I. Munro. 2002. Frequency estimation of internet packet streams with limited space. In *European Symposium on Algorithms*.
- [13] J.C. Deville and Y. Tillé. 1998. Unequal probability sampling without replacement through a splitting method. *Biometrika* 85, 1 (1998), 89–101.
- [14] Nick Duffield, Carsten Lund, and Mikkel Thorup. 2007. Priority sampling for estimation of arbitrary subset sums. *Journal of the ACM (JACM)* 54, 6 (2007), 32.
- [15] C. Estan, K. Keys, D. Moore, and G. Varghese. 2004. Building a better NetFlow. *ACM SIGCOMM Computer Communication Review* 34, 4 (2004), 245–256.
- [16] C. Estan and G. Varghese. 2003. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Transactions on Computer Systems (TOCS)* 21, 3 (2003), 270–313.
- [17] M. Ghashami, E. Liberty, and J.M. Phillips. 2016. Efficient Frequent Directions Algorithm for Sparse Matrices. *KDD* (2016).
- [18] P. B. Gibbons and Y. Matias. 1998. New sampling-based summary statistics for improving approximate query answers. *ACM SIGMOD Record* 27, 2 (1998), 331–342.
- [19] X. He, J. Pan, O. Jin, T. Xu, B. Liu, T. Xu, Y. Shi, A. Atallah, R. Herbrich, S. Bowers, and others. 2014. Practical lessons from predicting clicks on ads at facebook. In *International Workshop on Data Mining for Online Advertising*.
- [20] D. Hillard, S. Schroedl, E. Manavoglu, H. Raghavan, and C. Leggetter. 2010. Improving ad relevance in sponsored search. In *WSDM*. ACM, 361–370.
- [21] R. M. Karp, S. Shenker, and C. H Papadimitriou. 2003. A simple algorithm for finding frequent elements in streams and bags. *ACM Transactions on Database Systems (TODS)* 28, 1 (2003), 51–55.
- [22] E. Liberty. 2013. Simple and deterministic matrix sketching. In *KDD*.
- [23] G. Manku and R. Motwani. 2002. Approximate frequency counts over data streams. In *VLDB*.
- [24] A. Metwally, D. Agrawal, and A. El Abbadi. 2005. Efficient computation of frequent and top-k elements in data streams. In *ICDT*.
- [25] J. Misra and D. Gries. 1982. Finding repeated elements. *Science of computer programming* 2, 2 (1982), 143–152.
- [26] M. Mitzenmacher, T. Steinke, and J. Thaler. 2012. Hierarchical heavy hitters with the space saving algorithm. In *Meeting on Algorithm Engineering & Experiments*. 160–174.
- [27] M. Richardson, E. Dominowska, and R. Ragno. 2007. Predicting clicks: estimating the click-through rate for new ads. In *WWW*. ACM, 521–530.
- [28] V. Sekar, N. Duffield, O. Spatscheck, J. van der Merwe, and H. Zhang. 2006. LADS: large-scale automated DDOS detection system. In *USENIX*.
- [29] A. Shrivastava, A. C. König, and M. Bilenko. 2016. Time Adaptive Sketches (Ada-Sketches) for Summarizing Data Streams. *SIGMOD* (2016).
- [30] Mario Szegedy. 2006. The DLT priority sampling is essentially optimal. In *STOC*.
- [31] D. Vengerov, A.C. Menck, M. Zait, and S.P. Chakkappen. 2015. Join size estimation subject to filter conditions. *VLDB* (2015).
- [32] Y. Zhang, S. Singh, S. Sen, N. Duffield, and C. Lund. 2004. Online identification of hierarchical heavy hitters: algorithms, evaluation, and applications. In *Internet Measurement Conference (IMC)*.