

Higher-Fidelity Frugal and Accurate Quantile Estimation Using a Novel Incremental *Discretized* Paradigm

Anis Yazidi, *Member, IEEE*, Hugo Hammer, and B. John Oommen, *Fellow, IEEE*

Abstract—Traditional pattern classification works with the moments of the distributions of the features and involves the estimation of the means and variances etc. As opposed to this, more recently, research has indicated the power of using the Quantiles of the distributions because they are more robust and applicable for non-parametric methods. The estimation of the quantiles is even more pertinent when one is mining data streams. However, the complexity of quantile estimation is much higher than the corresponding estimation of the mean and variance, and this increased complexity is more relevant as the size of the data increases. Clearly, in the context of “infinite” data streams, a computational and space complexity that is linear in the size of the data is definitely not affordable. In order to alleviate the problem complexity, recently, a very limited number of studies have devised *incremental* quantile estimators [1], [2]. Estimators within this class resort to updating the quantile estimates based on the most recent observation(s), and this yields updating schemes with a very small computational footprint – a constant-time (i.e., $O(1)$) complexity. In this article, we pursue this research direction and present an estimator that we refer to as a Higher-Fidelity Frugal [1] quantile estimator. Firstly, it guarantees a substantial advancement of the family of Frugal estimators introduced in [1]. The highlight of the present scheme is that it works in the discretized space, and it is thus a pioneering algorithm within the theory of discretized algorithms¹. The convergence results that we have proven are based on the theory of Stochastic Point Location (SPL) [3], which we advocate as a new tool for solving a large class of online estimation problems. Extensive simulation results show that our estimator outperforms the original Frugal algorithm in terms of both speed and accuracy.

Index Terms—Quantile estimation, Stochastic Point Location, Discretized Estimation.

The third author is grateful for the partial support provided by NSERC, the Natural Sciences and Engineering Research Council of Canada. A very preliminary version of this work appeared in ADMA 2017, the 2017 *International Conference on Advanced Data Mining and Applications*, held in Singapore, November 5-6, 2017.

A. Yazidi and H. Hammer are with the Department of Computer Science, OsloMet Oslo Metropolitan University, Oslo, Norway.

B. J. Oommen is a *Chancellor’s Professor*, a *Fellow: IEEE* and a *Fellow: IAPR*. The Author also holds an *Adjunct Professorship* with the Dept. of ICT, University of Agder, Norway. His work was partially supported by NSERC, the Natural Sciences and Engineering Research Council of Canada.

¹The fact that discretized Learning Automata schemes are superior to their continuous counterparts has been clearly demonstrated in the literature. This is the first paper, to our knowledge, that proves the advantages of discretization within the domain of quantile estimation.

I. INTRODUCTION

Estimation is probably the most fundamental and central problem in many areas of engineering and computer science. The entire training phase of classification deals with estimation in one way or the other. While solutions to estimating the mean (and central or non-central moments) of a distribution have been well established for centuries, we consider the problem of estimating the quantiles of a distribution with minimal time and space requirements.

Apart from the phenomenon of estimation, there are three rather distinct computational paradigms that have emerged within the general area of computational intelligence listed below:

- 1) The first of these involves the Stochastic Point Location SPL problem [3] where the Learning Mechanism (LM) attempts to learn a point on the “line” when all that it receives are signals from a random environment, i.e., whether *it* is to the “Left” or “Right” of the unknown point. This point that the LM attempts to learn may be, for example, a parameter of a control system. This problem has been briefly addressed in Section I-A.
- 2) The second of these involves the concept of discretization. Unlike learning in a continuous probability space, it has been shown that in the field of Learning Automata (LA), it is advantageous to discretize the probability space. Discretized LA are, generally speaking, both faster and more accurate than their corresponding continuous counterparts. A brief overview of discretization has been included in Section I-B.
- 3) The third of these are the unique issues encountered when one seeks to estimate the quantiles of a distribution rather than the mean or central/non-central moments of a distribution in an *incremental* manner. A survey of estimation methods for the estimation of quantiles is given in Section I-C.

Conceptually, the fundamental contribution of this paper is to present a single solution that represents the confluence of these three distinct paradigms. By first presenting these three paradigms in the next three subsections, we will also be surveying the state-of-the-art in the respective fields.

A. The SPL and its Solutions

To place our contributions in the right perspective, we briefly review² the state of the art of the SPL problem, whose formulation and solution is central to our approach. The SPL problem, in its most elementary formulation, assumes that there is a Learning Mechanism (LM) whose task is to determine the optimal value of some variable (or parameter), x . We assume that there is an optimal choice for x – an unknown value, say $x^* \in [0, 1)$. The SPL involves inferring the value x^* . Although the mechanism does not know the value of x^* , the SPL assumes that it has responses from an intelligent “Environment” (synonymously, referred to as the “Oracle”), Ξ , that is capable of informing it whether any value of x is too small or too big. To render the problem both meaningful and distinct from its deterministic version, we would like to emphasize that the response from this Environment is assumed “faulty.” Thus, Ξ may tell us to increase x when it should be decreased, and *vice versa*. However, to render the problem tangible, in [3] the probability of receiving an intelligent response was assumed to be $p > 0.5$, in which case Ξ was said to be *Informative*. Note that the quantity “ p ” reflects on the “effectiveness” of the Environment. Thus, whenever the current $x < x^*$, the Environment correctly suggests that we increase x with probability p . It simultaneously could have incorrectly recommended that we decrease x with probability $(1 - p)$. The converse is true for $x \geq x^*$.

The first known solution to the problem is due to Oommen [3], who pioneered the study of the SPL when he proposed and analyzed an algorithm that operates on a discretized search space while interacting with an informative Environment (i.e., $p > 0.5$). The search space was first sliced into N sub-intervals at the positions $\{0, \frac{1}{N}, \frac{2}{N}, \dots, \frac{N-1}{N}, 1\}$, where a larger value of N ultimately implied a more accurate convergence to x^* . The algorithm then did a controlled random walk on this space by “obediently” following the Environment’s advice in the discretized space. In spite of the Oracle’s erroneous feedback, this discretized solution was proven to be ϵ -optimal.

An novel alternate *parallel* strategy that combined LA and pruning was used in [4] to solve the SPL. By utilizing the response from the environment, the authors of [4] partitioned the interval of search into three disjoint subintervals, eliminating at least one of the subintervals from further search, and by recursively searching the remaining interval(s) until the search interval was at least as small as the required resolution. In a subsequent work [5], Oommen *et al.* introduced the Continuous Point Location with Adaptive d-ARY Search (CPL-AdS), which was a generalization of the work in [4].

An extension of the latter work to the case of dynamic environments was reported in [6]. Recently Yazidi *et al.* [7] proposed a *hierarchical* searching scheme for solving

the SPL problem. This solution involved partitioning the line in a hierarchical tree-like manner, and then moving to relatively distant points, as characterized by those along the paths of the tree.

B. The Phenomenon of Discretization

We now present a very brief overview of the second phenomenon alluded to above, namely that of discretization. Historically, the concept of discretizing the probability space was pioneered by Thathachar and Oommen in their study on Reward-Inaction LA [8], and since then, it has catalyzed a significant research in the design of discretized LA [9]–[13]. In these algorithms, the probability changes are made in jumps and not continuously, and thus, the speed of the learning process is increased, especially as one approaches the optimal solution. Discretization is also beneficial when it concerns issues related to implementation and representation. Since such algorithms use integer (as opposed to real number) representations, they permit addition (as opposed to multiplication) operations. Some of the existing results about discretized automata are found in [8], [9], [12]–[16]. Indeed, the fastest reported LAs are the discretized pursuit and estimator algorithms [9], [15], [16]. Recently, there has been an upsurge of research interest in solving resource allocation problems based on novel discretized LA [10], [11], in which the authors proposed a solution to the class of *Stochastic Nonlinear Fractional Knapsack* problems where resources had to be allocated based on incomplete and noisy information. The latter solution was thereafter applied to resolve the web-polling problem, and to the problem of determining the optimal size required for training a classifier. By virtue of discretization, the estimator that we propose realizes fast adjustments of the running estimates by performing “jumps”, and it is thus able to robustly and quickly track changes in the parameters of the distribution after a switch has occurred in the environment.

C. Quantile Estimation

The estimation of the quantiles of a distribution is far more complex than the estimation of its mean or central/non-central moments. Indeed, unlike the mean or central/non-central moments, the occurrence of a few samples can drastically change the location of many (if not all) the quantiles.

Historically, it is worth mentioning that the seminal paper of Robbins and Monro [17] which established the field of research called “stochastic approximation” [18], had already, decades ago, included an incremental quantile estimator as a proof-of-concept application for the vast arena for the theory of stochastic approximation. An extension of the latter quantile estimator, which first appeared as an example in [17], was further developed in [19] in order to handle the case of “extreme quantiles”. It should be mentioned that the estimator provided by Tierney in [20] falls under the same umbrella of the

²This review can be abridged or even deleted if requested by the Referees.

example provided in [17]; it can thus be seen as an extension of it.

The application domains for the use of quantiles, especially in the recently-introduced “Anti”-Bayesian methods of classification, require the robust computation of the quantiles of the distributions encountered. Further, in many network monitoring applications, quantiles are key indicators for monitoring the performance of the system. For instance, system administrators are interested in monitoring the 95% response time of a web-server so that they can constrain it to be under a certain threshold. Quantile tracking is also useful for detecting abnormal events, and in intrusion detection systems, in general. However, the immense traffic volume of high speed networks impose some computational challenges, namely, that the storage is limited, and the requirement on the computation done on the data is that it needs to be achieved in a “one pass” manner.

A body of research has been focused on quantile estimation from data streams without making any specific assumption on the distribution of the data samples. But since we are working on data streams, it is prudent to review some of the related work that concentrates on estimating quantiles from data streams³. The most representative work of this type of “streaming” quantile estimator is due to the seminal work of Munro and Paterson [21]. In [21], these authors described a p -pass algorithm for selection using $O(n^{1/(2p)})$ space for any $p \geq 2$. Cormode and Muthukrishnan [22] proposed a more space-efficient data structure, called the Count-Min sketch, which was inspired by Bloom filters, where one estimates the quantiles of a stream as the quantiles of a random sample of the input. The key idea here was to maintain a random sample of an appropriate size to estimate the quantile, where the premise was to select a subset of elements whose quantile approximated the true quantile. From this perspective, the latter body of research required a certain amount of memory that increased as the required accuracy of the estimator increases [23]. Examples of these works are [21], [23]–[26]. Guha and McGregor [26] advocated the use of random-order data models in contrast to adversarial-order models. They showed that computing the median requires an exponential number of passes in an adversarial model, while it required $O(\log \log n)$ in a random-order model.

In [27], the authors proposed a modification of the stochastic approximation algorithm [20] in order to allow an update similar to the well-known “Exponentially Weighted Moving Averages” form for updates. This modification is particularly helpful in the case of non-stationary environments in order to cope with non-stationary data. Thus, the quantile estimate is a weighted combination of the new data that has arrived and the previously-computed estimate. Indeed, a “weighted”

update scheme is applied to incrementally build local approximations of the distribution function in the neighborhood of the quantiles.

D. Contributions

As mentioned in the introductory section, this paper brings together the principles recorded in Sections I-A, I-B and I-C. It introduces a novel *discretized* quantile estimator based on the principles of SPL. Although we had earlier solved the binomial estimation problem using discretized estimators [28]–[30], this is the first solution to quantile estimation that involves the SPL-based solution. Our estimator, the Higher-Fidelity Frugal (H-FF), outperforms the Original Frugal (OF) [1], that is also discretized. From the above survey, the reader will observe that the entire field of utilizing the concepts of discretization in quantile estimation has been under-investigated, which is, precisely, our primary contributions. Thus, apart from the above, we can catalogue the contributions of the paper as follows:

- While the OF algorithm [1] fails to generalize the estimation rules for the median case, our H-FF design, which is based on the SPL, accommodates the case of the median as a simple instantiation.
- The experimental results that we report, demonstrate that our H-FF algorithm outperforms the state-of-the-art discretized OF algorithm in terms of both speed and accuracy. This is especially true in the case of estimating quantiles close to the median.
- Our H-FF follows the principles of SPL introduced in [31]. The main difference between our estimator and the original SPL is that there is a non zero probability that, in our present updating scheme, the estimate remains unchanged at the next time instant. This will be explained in greater detail in Section II-B.

II. ON ENHANCING THE FRUGAL ESTIMATOR

Since our contribution falls into the family of *Incremental* Quantile Estimators, we now present an overview of this class of estimators.

A. Incremental Quantile Estimators

An incremental estimator, by definition, resorts to the last observation(s) in order to update its estimate. The research on developing incremental quantile estimators is sparse. Probably, one of the outstanding early and unique examples of incremental quantile estimators is due to Tierney, proposed in 1983 [20], and which resorted to the theory of stochastic approximation. Applications of Tierney’s algorithm to network monitoring can be found in [32]. The shortcoming of Tierney estimator [20] is that it requires the incremental constructions of local approximations of the distribution function in the neighborhood of the quantiles, and this increases the

³As we will explain later, these related works require some memory restrictions which renders our work to be radically distinct from them. Our approach requires storing only a single sample value in order to update the estimate.

complexity of the algorithm. Our goal is to present an algorithm that does not involve any local approximations of the distribution function. Recently, a generalization of the Tierney’s [20] algorithm was proposed by the authors of [27], where the authors proposed a batch update of the quantile, where the quantile is updated every $M \geq 1$ observations.

In the same context of incremental estimators, Ma, Muthukrishnan and Sandler [1] recently devised an innovative incremental quantile estimator⁴ called the Frugal scheme, that follows randomized rules of updates. The first algorithm presented in the manuscript of Ma, Muthukrishnan and Sandler [1] is a Frugal approach for estimating the median. The procedure for estimating the median is simple but also “surprising”: One increments the estimate of the median by a fixed amount Δ ($\Delta > 0$) whenever the observation from the data stream is larger than the median, and decrements the estimate of the median by Δ whenever the observation is smaller than the corresponding estimate. Nevertheless, the Frugal algorithm presented later in the same manuscript in order to tackle any quantile estimate (apart from the median), is not a generalization of the median case. In fact, according to the general update equations, if we are attempting to find the 50% quantile (median) of the data stream, we need to increment up randomly with 50% probability (for observations larger than the median estimate) and decrement down randomly with 50% probability (for observations smaller than the median estimate). Thus, intuitively, the Frugal [1] algorithm fails to generalize the median case as we observe that the randomization is unnecessary for estimating the median. Moreover, we can intuitively infer that the Frugal algorithm will suffer also from the “unnecessary” randomization for quantile estimates that fall in neighborhood of 50%.

In [2], Yazidi and Hammer devised a truly multiplicative incremental quantile estimation algorithm. The main difference between that and the current work is that the latter algorithm operates on a continuous space, while this present work is in a discretized space.

When it comes to memory efficient methods that require a small storage footprint, histogram based methods form an important class. Viewed from this perspective, a representative work is due to Schmeiser and Deutsch [34] who proposed the use of equidistant bins, where the boundaries are adjusted online. Arandjelovic *et al.* [35] used a different idea than equidistant bins by attempting to maintain bins in a manner that maximizes the entropy of the corresponding estimate of the historical data distribution, and where the bin boundaries were adjusted in an online manner.

In [36], Jain *et al.* resorted to five markers so as to track the quantile, where the markers corresponded to different quantiles and the min and max of the observations. Their concept was similar to the notion of histograms,

where each marker had two measurements, its height and its position. By definition, each marker had some ideal position, and some adjustments were made so as to keep it in its ideal position by counting the number of samples that exceeded the marker. Thus, for example, if the marker corresponded to the 80% quantile, its ideal position would be around the point corresponding to the 80% of the data points below the marker. Subsequently, based on the positions of the markers, the quantiles were computed by modeling it such that the curve passing through three adjacent markers was parabolic, and by using piecewise parabolic prediction functions⁵.

In [37], the authors proposed a memory-efficient method, based on the algorithm from [36], for the simultaneous estimation of several quantiles using interpolation methods and a grid structure, where each internal grid point was updated upon receiving an observation. Their approximation relied on using linear and parabolic interpolations, while the tails of the distribution were approximated using exponential curves.

A notable work dealing with the simultaneous estimation of the quantiles using elements from the theory of stochastic approximation is due to Cao *et al.* [38]. Here, the authors resorted to interpolation by defining a distance measure between the interpolated quantiles so as to ensure that there were no “crossings” between the monotonic quantile estimates. Nevertheless, the interpolation used estimates of the “density” as in [27] and [20], which is an operation that increases the complexity.

Finally, it is worth mentioning that an important research direction that has received little attention in the literature revolves around updating the quantile estimates under the assumption that portions of the data are deleted. Such an assumption is realistic in many real-life settings where data needs to be deleted due to the occurrence of errors, or because they are out-of-date and thus should be replaced. The deletion triggered a recomputation of the quantile [38], which is considered a complex operation. The case of deleted data is more challenging than the case of insertion of new data, because data insertion can be handled easily using either sequential or batch updates, while quantile update upon deletion requires more complex update operations.

B. The Higher-Fidelity Frugal Estimator

To motivate our work, we concur with Arandjelovic *et al.* [35] who remark that most quantile estimation algorithms are not single-pass algorithms and are, thus, not applicable for streaming data. On the other hand, the single pass algorithms are concerned with the exact computation of the quantile and thus require a storage space of the order of the size of the data which is clearly an unfeasible condition in the context of “Big Data” streams. Thus, the work on quantile estimation using

⁴With some insight, one sees that this elegant median estimation procedure is similar to the Boyer and Moore algorithm [33] for computing the majority item in a stream, using only a single pass.

⁵Clearly, though, such an approach would not be able to handle the case of non-stationary quantile estimation as the positions of the markers would be affected by stale data points.

more than one pass, or storage of the same order of the size of the observations seen so far, is not relevant in the context of this paper. We also affirm the need for storage-constrained and single-pass algorithms.

In this article, we extend the results from Frugal [1] and present a Higher-Fidelity Frugal (H-FF) scheme where the median can be seen as an instantiation of our algorithm and not as exceptional case that requires a different set of rules. In addition, our H-FF scheme is shown to be faster and more accurate than the original Frugal scheme [1]. For the rest of the paper, in order to avoid confusion, we will refer to the original Frugal algorithm due to Ma, Muthukrishnan and Sandler [1], as the Original Frugal (OF). As mentioned earlier, our H-FF algorithm is based on the theory of Stochastic Point Location [3], and although the latter theory has found applications within discretized binomial and multinomial estimation in [29], as we shall see, its application here is unique. In addition, one can observe that the binomial/multinomial discretized estimators proposed by Yazidi *et al.* in [28], [30] and Frugal [1] are similar. In fact, if we use the same update equations as in [28], [30] with the “binary” observation being whether the current estimate sample is larger than the current estimate, then, interestingly, we obtain the OF scheme [1]!

Let $Q_i = a + i \cdot \frac{(b-a)}{N}$ and suppose that we are estimating the quantile in the interval⁶ $[a, b]$. Note $Q_0 = a$ and $Q_N = b$. Let Δ be $\frac{(b-a)}{N}$. Further, we suppose that the estimate at each time instant $\hat{Q}(n)$ takes values from the $N + 1$ possible values, i.e., $Q_i = a + i \cdot \Delta$, where $0 \leq i \leq N$.

For the sake of completeness, we will give the update equations for the OF algorithm introduced in [1]. Let $x(n)$ denote the observation at time instant n . Please note that the equations are slightly modified so as to obtain estimates within $[a, b]$. In addition, the step size Δ has a general form and is not limited to unity as done in [1].

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \text{Min}(\hat{Q}(n) + \Delta, b), \\ &\quad \text{If } \hat{Q}(n) \leq x(n) \text{ and } \text{rand}() \leq q, \end{aligned} \quad (1)$$

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \text{Max}(\hat{Q}(n) - \Delta, a), \\ &\quad \text{If } \hat{Q}(n) > x(n) \text{ and } \text{rand}() \leq 1 - q, \end{aligned} \quad (2)$$

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \hat{Q}(n), \\ &\quad \text{Otherwise,} \end{aligned} \quad (3)$$

where $\text{Max}(\cdot, \cdot)$ and $\text{Min}(\cdot, \cdot)$ denote the max and min operator of two real numbers while $\text{rand}()$ is a random number generated uniformly in $[0, 1]$.

Our H-FF algorithm has two different update equations depending on whether the quantile we are estimating is larger or smaller than the median.

Update equation for $q \leq 0.5$:

⁶Throughout this paper, there is an implicit assumption that the true quantile lies in $[a, b]$. However, this is not a limitation of our scheme; the proof is valid for any bounded and probably non-bounded function.

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \text{Min}(\hat{Q}(n) + \Delta, b), \\ &\quad \text{If } \hat{Q}(n) \leq x(n) \text{ and } \text{rand}() \leq \frac{q}{1-q}, \end{aligned} \quad (4)$$

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \text{Max}(\hat{Q}(n) - \Delta, a), \\ &\quad \text{If } \hat{Q}(n) > x(n), \end{aligned} \quad (5)$$

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \hat{Q}(n), \\ &\quad \text{Otherwise.} \end{aligned} \quad (6)$$

Update equations for $q > 0.5$:

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \text{Min}(\hat{Q}(n) + \Delta, b), \\ &\quad \text{If } \hat{Q}(n) \leq x(n), \end{aligned} \quad (7)$$

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \text{Max}(\hat{Q}(n) - \Delta, a), \\ &\quad \text{If } \hat{Q}(n) > x(n) \text{ and } \text{rand}() \leq \frac{1-q}{q}, \end{aligned} \quad (8)$$

$$\begin{aligned} \hat{Q}(n+1) &\leftarrow \hat{Q}(n), \\ &\quad \text{Otherwise.} \end{aligned} \quad (9)$$

The proof of the properties of the updating scheme follows.

Theorem 1. *Let us assume that we are estimating the q -th quantile of the distribution, i.e., $Q^* = F_X^{-1}(q)$. Then, applying the updating rules given by Equations (4) - (6) for the case when $q \leq 0.5$, and Equations (7) - (9) when $q > 0.5$ yields: $\lim_{N \rightarrow \infty} \lim_{n \rightarrow \infty} E(\hat{Q}(n)) = Q^*$.*

Proof. We shall prove the above by analyzing the properties of the underlying Markov chain, which is specified by the rules given by Equations (4) - (6) for the case where $q \leq 0.5$ and Equations (7) - (9) in case where $q > 0.5$. The states of the chain are the integers $\{0, 1, 2, \dots, N\}$, and these correspond to the values $\{Q_0, Q_1, Q_2, \dots, Q_N\}$, respectively.

By considering Equations (4) - (6), we deduce the state transition probabilities for $q \leq 0.5$ as:

$$h_{i,i+1} = \frac{q}{1-q}(1 - F_X(Q_i)), 0 \leq i \leq N-1, \quad (10)$$

$$h_{i,i-1} = F_X(Q_i), 1 \leq i \leq N, \quad (11)$$

$$h_{i,i} = 1 - h_{i,i-1} - h_{i,i+1}, 0 < i < N. \quad (12)$$

This is due to the fact that the probability that $\hat{Q}(n) \leq x(n)$ can be expressed in terms of the CDF function $F_X(\cdot)$ namely, that $\text{Prob}(\hat{Q}(n) \leq x(n)) = 1 - F_X(\hat{Q}(n))$.

The transitions for the boundary states (when $q \leq 0.5$) are self-loops, and obey:

$$h_{0,0} = 1 - \frac{q}{1-q}(1 - F_X(Q_0)), \quad (13)$$

$$h_{N,N} = \frac{q}{1-q}(1 - F_X(Q_N)). \quad (14)$$

Similarly, by observing Equations (7) - (9), the state

transition probabilities for $q > 0.5$ are:

$$h_{i,i+1} = 1 - F_X(Q_i), 0 \leq i \leq N-1, \quad (15)$$

$$h_{i,i-1} = \frac{1-q}{q} F_X(Q_i), 1 \leq i \leq N, \quad (16)$$

$$h_{i,i} = 1 - h_{i,i-1} - h_{i,i+1}, 0 < i < N, \quad (17)$$

and the transitions for the boundary states (when $q > 0.5$) are self-loops, and obey:

$$h_{0,0} = \frac{1-q}{q} F_X(Q_0), \quad (18)$$

$$h_{N,N} = 1 - \frac{1-q}{q} (1 - F_X(Q_N)). \quad (19)$$

We shall now compute π_i the stationary (or equilibrium) probability of the chain being in state i . Clearly H represents a single closed communicating class whose periodicity is unity. The chain is ergodic, and the limiting probability vector is given by the eigenvector of H^T corresponding to the eigenvalue unity. The vector of steady state probabilities $\Pi = [\pi_0, \pi_1, \dots, \pi_N]^T$ can be computed using $H^T \Pi = \Pi$.

Whether we are applying the updating Equations (4) - (6) when $q \leq 0.5$, or Equations (7) - (9) when $q > 0.5$, the rules obey the Markov chain with transition matrix $H = [h_{ij}]$, where:

Consider first the stationary probability of being in state 0, π_0 . Expanding the first row of Equation (20) yields:

$$\pi_0 h_{0,0} + \pi_1 h_{1,0} = \pi_0 \implies \pi_1 = \frac{(1-h_{0,0})}{h_{1,0}} = \frac{h_{0,1}}{h_{1,0}} \pi_0. \quad (21)$$

Expanding the second row of Equation (20) and substituting (21) yields:

$$\pi_0 h_{0,1} + \pi_2 h_{2,1} = \pi_1 \implies \pi_2 = \frac{h_{1,2}}{h_{2,1}} \pi_1. \quad (22)$$

Arguing in the same manner, and after some algebraic simplifications, we obtain the recurrence relation for $0 < k \leq N$:

$$\pi_{k-1} = \frac{h_{k,k-1}}{h_{k-1,k}} \pi_k, \quad (23)$$

which, on reversing the recursion, yields for $0 \leq k < N$

$$\pi_{k+1} = \frac{h_{k,k+1}}{h_{k+1,k}} \pi_k. \quad (24)$$

Let z be the index for which $z\Delta \leq Q^* < (z+1)\Delta$, where Q^* the true quantile to be estimated.

The crucial part of our proof is to reformulate Π in terms of π_z and π_{z+1} , using Equations (23) and (24). More specifically, for $j \in \{0, 1, \dots, z-1\}$ we have:

$$\pi_j = \pi_z \prod_{k=z}^{j+1} \frac{h_{k,k-1}}{h_{k-1,k}}. \quad (25)$$

Correspondingly, and arguing in an analogous manner, for $j \in \{z+2, \dots, N\}$ we have:

$$\pi_j = \pi_{z+1} \prod_{k=z+1}^{j-1} \frac{h_{k,k+1}}{h_{k+1,k}}. \quad (26)$$

In other words, we represent Π in terms of two of its components, namely, π_z and π_{z+1} . We are now ready to define the upper bound U for Π :

$$U[i, z] = \begin{cases} \pi_z M^{z-i} & \text{if } i \leq z \\ \pi_{z+1} M^{i-(z+1)} & \text{if } i \geq z+1, \end{cases}$$

where: $M = \max \left[\max_{i \leq z} \frac{h_{i,i-1}}{h_{i-1,i}}, \max_{i \geq z+1} \frac{h_{i,i+1}}{h_{i+1,i}} \right]$. As seen, the definition of M clearly makes U an upper bound for Π almost everywhere, except for $i = z+1$. Our final goal is to show that as the resolution N goes to infinity, U goes to zero outside the small interval $[z\Delta, (z+1)\Delta]$, implying that:

$$\lim_{N \rightarrow \infty} U[i, z] = 0, \text{ if } i \notin \{z, z+1\}$$

We shall argue that the latter is guaranteed to happen if we have $0 < \frac{h_{i,i-1}}{h_{i-1,i}} < 1$ for $i \in \{1, \dots, z\}$, and $0 < \frac{h_{i,i+1}}{h_{i+1,i}} < 1$ for $i \in \{z+1, \dots, N-1\}$. This is because when these conditions are met, we obtain $0 < M < 1$. We argue this by considering the equilibrium (asymptotic) value of $E(\Pi(n))$ for any finite N . This argument can be separated into three different cases as in [3]:

- 1) The first case is when $z\Delta$ is close to a . In this case, the maximum is quickly reached and then geometrically falls away.
- 2) When $z\Delta$ is close to b , the value of π_i geometrically increases but when the maximum is reached, it quickly falls away. For both these cases when $N \rightarrow \infty$, most of the probability mass will be centered in the small interval $[z\Delta, (z+1)\Delta]$.
- 3) The third case is slightly more complex because it involves $z\Delta$ being away from either end. This case must be broken down into two distinct geometric series, one representing the geometric series from π_0 to π_z and the other from π_{z+1} to π_N . The first series increases until it reaches the maximum at π_z . The increase is geometric (or rather, exponential as $N \rightarrow \infty$), and the geometric ratio is bounded by the bound given by the quantity M above. The second series starts at the maximum at the value π_{z+1} and then decreases until π_N is reached. Again, the decrease is geometric (i.e., exponential as $N \rightarrow \infty$), and the geometric ratio is bounded by the quantity M above. In this case, the probability mass will be centered within the small interval $[z\Delta, (z+1)\Delta]$ as $N \rightarrow \infty$ because of the law of the sum of the elements of a geometric series possessing a common ratio which is greater than unity.

We shall now demonstrate that $\frac{h_{i-1,i}}{h_{i,i-1}} < 1$ for $i \in \{1, \dots, z\}$ and that $\frac{h_{i,i+1}}{h_{i+1,i}} < 1$ for $i \in \{z+1, \dots, N-1\}$. This reduces to demonstrating that $\frac{h_{i-1,i}}{h_{i,i-1}} < 1$ for $i \in \{1, \dots, z\}$ and that $\frac{h_{i-1,i}}{h_{i,i-1}} > 1$ for $i \in \{z+2, \dots, N\}$, implying that:

$$\frac{h_{i-1,i}}{h_{i,i-1}} = \frac{q(1 - F_X(Q_{i-1}))}{(1-q)F_X(Q_i)}.$$

$$\begin{bmatrix}
h_{0,0} & h_{0,1} & 0 & \cdot & \cdot & \cdot & \cdot & 0 \\
h_{1,0} & h_{1,1} & h_{1,2} & 0 & \cdot & \cdot & \cdot & 0 \\
0 & h_{2,1} & h_{2,2} & h_{2,3} & 0 & \cdot & \cdot & 0 \\
\vdots & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\
\cdot & \cdot & \cdot & h_{k,k-1} & 0 & h_{k,k+1} & \cdot & \cdot \\
\vdots & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\
\vdots & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\
0 & \cdot & \cdot & \cdot & \cdot & h_{N-1,N-2} & h_{N-1,N-1} & h_{N-1,N} \\
& & & & & 0 & h_{N,N-1} & h_{N,N}
\end{bmatrix}^T
=
\begin{bmatrix}
\pi_0 \\
\pi_1 \\
\pi_2 \\
\vdots \\
\pi_{N-1} \\
\pi_N
\end{bmatrix}
=
\begin{bmatrix}
\pi_0 \\
\pi_1 \\
\pi_2 \\
\vdots \\
\pi_{N-1} \\
\pi_N
\end{bmatrix}
\quad (20)$$

Let us consider the difference between the numerator and the denominator as:

$$\begin{aligned}
q(1 - F_X(Q_{i-1})) - (1 - q)F_X(Q_i) &> \\
q(1 - F_X(Q_i)) - (1 - q)F_X(Q_i) & \quad (27) \\
&= q - F_X(Q_i).
\end{aligned}$$

where, in Equation (27), we have resorted to the fact that $F_X(Q_{i-1}) < F_X(Q_i)$.

Similarly, we obtain the following inequality by using fact that $F_X(Q_{i-1}) < F_X(Q_i)$:

$$\begin{aligned}
q(1 - F_X(Q_{i-1})) - (1 - q)F_X(Q_i) &< \\
q(1 - F_X(Q_{i-1})) - (1 - q)F_X(Q_{i-1}), & \quad (28) \\
&= q - F_X(Q_{i-1}).
\end{aligned}$$

Using the two above inequalities, i.e., Equations (27) and (28), we conclude that

$$q - F_X(Q_i) < q(1 - F_X(Q_{i-1})) - (1 - q)F_X(Q_i) < q - F_X(Q_{i-1}). \quad (29)$$

We know that for $i \in \{1, \dots, z\}$, $Q_i < Q^*$ is true, and consequently $F_X(Q_i) < F_X(Q^*) = q$ by virtue of the monotonicity of the CDF function. Using the inequality (29), we can thus conclude that $\frac{h_{i,i-1}}{h_{i-1,i}} < 1$ for $i \in \{1, \dots, z\}$.

Similarly, we can prove the second case. Indeed, we know that for $i \in \{z+2, \dots, N\}$, the inequality $Q_{i-1} > Q^*$ is true, and consequently $F_X(Q_{i-1}) > F_X(Q^*) = q$ by virtue of the monotonicity of the CDF function. Therefore, in this case, $q - F_X(Q_{i-1}) < 0$, and consequently, $\frac{h_{i,i-1}}{h_{i-1,i}} > 1$ for $i \in \{z+2, \dots, N\}$. Hence the theorem. \square

C. Salient Differences between the H-FF, SPL and OF

It is pertinent to mention that there are some fundamental differences between the H-FF and the SPL, both with regard to their *computational paradigms* and with regard to their respective *analyses*. There are also some fundamental differences between the H-FF and the OF schemes. We state them briefly below.

1) *Differences between the Paradigms of the H-FF and SPL:* The following are the differences between the *paradigms* of the H-FF and SPL:

- Although the rationale for updating in the H-FF is *apparently* similar to that of the SPL algorithm [3], there are some fundamental differences. First,

we emphasize that the SPL has a significant advantage. Indeed, the SPL assumes the existence of an ‘‘Oracle’’, the presence of which is, unarguably, a ‘‘bonus’’. In our case, since there is no ‘‘Oracle’’, the H-FF scheme has to simulate such an entity. Or more precisely, it has to infer the behavior of a fictitious ‘‘Oracle’’ from the incoming samples.

- Further, unlike the SPL, the H-FF has no specific LM either. The learning properties of the LM must now be encapsulated into the estimation procedure.

2) *Differences between the Analyses of the H-FF and SPL:* The following are the differences between the *analyses* of the H-FF and SPL:

- From a cursory perspective, it could *appear* as if the Markov Chain that we have presented, and its analysis, are rather identical to those presented in [3]. However, although the similarities are few, the differences are more vital. The main differences are the following:

- 1) First of all, unlike the original SPL, there is a non zero probability that in our present updating scheme, the estimate remains unchanged at the next time instant.
- 2) As opposed to original SPL, in our case, the scheme never stays at the same state at the next time instant, except at the end states. Rather, the environment (our simulated ‘‘Oracle’’) directs the simulated LM to move to the right or to the left, or to stay at the same position.

- Unlike the work of [3], the probability that the ‘‘Oracle’’ suggests the move in the correct direction, is not constant over the states of the estimator’s state space. This is quite a significant difference, since it renders our model to be characterized by a Markov Chain with state-*dependent* transition probabilities.
- A major advantage of this estimator and SPL-based estimators, in general, is that they are, by design, adequate to dynamic environments. In fact, the estimator is memory-less, and this is a consequence of the Markovian property. Thus, whenever a change takes place in the unknown underlying value of the target quantile to be tracked, our H-FF will *instantly* change its search direction since the properties of transition probabilities of the underlying random walk, change too.

3) *Other Salient Differences between the H-FF and OF:*

- Our H-FF is “semi-randomized” in the sense that only one direction of the updates is randomized and not both directions as in the case of the OF algorithm. In fact, whenever $q \leq 0.5$, we observe that the randomization is only applied for moving to the left (decrementing the estimate with probability $\frac{q}{1-q}$ which is less than unity). Similarly, when estimating a quantile q such that $q > 0.5$, the randomization is only applied for moving to the right (incrementing the estimate with probability $\frac{1-q}{q}$, which is again strictly less than unity).
- A fundamental observation is that for the median case, i.e., when $q = 0.5$, we obtain the Frugal update proposed as an exceptional case that deviates from the main scheme in [1] since $\frac{q}{1-q} = 1$. Formally, the median is estimated as follows:

$$\widehat{Q}(n+1) \leftarrow \begin{aligned} & \text{Min}(\widehat{Q}(n) + \Delta, b) \\ & \text{if } \widehat{Q}(n) \leq x(n), \end{aligned} \quad (30)$$

$$\widehat{Q}(n+1) \leftarrow \begin{aligned} & \text{Max}(\widehat{Q}(n) - \Delta, a) \\ & \text{if } \widehat{Q}(n) > x(n). \end{aligned} \quad (31)$$

III. EXPERIMENTAL RESULTS

In order to demonstrate the strength of our scheme (denoted as H-FF), we have rigorously tested it and compared it to the OF estimator proposed in [1] for different distributions, under different resolution parameters, and in both dynamic and stationary environments. The results we have obtained are conclusive and demonstrate that the convergence of the algorithms conforms to the theoretical results, and proves the superiority of our design to the OF algorithm [1]. To do this, we have used data originating from different distributions, namely:

- Uniform in $[0, 1]$,
- Normal $N(0, 1)$,
- Exponential distribution with mean 1 and variance 1, and
- Chi-square distribution with mean 1 and variance 2.

In all the experiments, we chose a to be -8 and b to be 8 . Note that whenever the resolution was N , the estimate was moving with either an additive or subtractive step size equal to $\frac{b-a}{N}$. Thus, a larger value of the resolution parameter, N , implied a smaller step size, while a lower value of the resolution parameter, N , led to a larger step size. Initially, at time 0, the estimates were set to the value $Q_{\lfloor \frac{N}{2} \rfloor}$. The reader should also note that an additional aim of the experiments was to demonstrate the H-FF’s salient properties as a novel quantile estimator using only *finite* memory.

A. Comparison in Stationary Environments for Different Distributions

In this set of experiments, we examined various stationary environments. We used different resolutions, and as mentioned previously, we set $[a, b] = [-8, 8]$. In each

case, we ran an ensemble of 1,000 experiments, each consisting of 500 iterations.

In Tables I, II, III and IV, we report the estimation error for the OF and H-FF for different values of the resolutions, N , for the Uniform, Normal, Exponential and Chi-squared distributions respectively. We catalogue the results for different values of the quantile being estimated, namely, q : 0.1, 0.3 0.499, 0.7 and 0.9. From these tables we observe that the H-FF outperformed the OF in almost all the cases, i.e., for different distributions and for different resolutions. A general observation is that the error for both schemes diminished as we increased the resolution. For example, from Table I, we see that the error for $q = 0.1$ decreased from 0.144 to 0.044 as the resolution increased from 50 to 500.

A very intriguing characteristic of our estimator is that as the resolution increased, the estimation error diminished (asymptotically). In fact, the limited memory of the estimator did not permit us to achieve zero error, i.e., 100% accuracy. As noted in the theoretical results, the convergence centred around the smallest interval $[z\Delta, (z+1)\Delta]$ containing the true quantile. Informally speaking, a higher resolution increased the accuracy while a low resolution decreased the accuracy.

Another interesting remark is that both the OF and H-FF seemed to perform almost equally well for extreme quantiles, i.e., quantiles that are close to 0 or close to 1. However, as the true value of the quantile to be estimated became closer to 0.5, i.e, median, the H-FF had a markedly clearer superiority when compared to the OF.

The reader should note that the choice of 0.499 instead of 0.5 was deliberate in order to “avoid” using the exceptional rules presented with regard to the OF in [1], and that coincide with the rules of H-FF for the median. Thus, the estimation of the quantile for the value 0.499 was performed using the OF rules as per Equations (1) - (3) to avoid the unnecessary randomization of the OF around the median that could lead to higher errors, which was the earlier-mentioned shortcoming of the OF scheme.

Please note too that for the target values of the quantiles that were close to the initial point 0, the error was smaller than for those that are far away from initial point. Thus, for example, in Table I, the error was lowest for the 10% quantile which is 0.1, which in this case, is closer to 0 than any other quantile in the the table, namely, 0.3 0.499, 0.7 and 0.9.

Figure 1 depicts the case of estimating the 30% quantile for the four different distributions: the Uniform, Normal, Exponential and Chi-square. Similarly Figure 2, 3, 4 and 5 depict the case of estimating the 46%, 53%, 85% and 97% quantile respectively for the four different distributions: the Uniform, Normal, Exponential and Chi-square. In all case, the resolution set as 150.

From our experiments, we observed that our H-FF algorithm approached the true value for all the four distributions asymptotically over time. However, when

TABLE I: The estimation error for the OF and H-FF algorithms for the Uniform distribution and for different values of the resolutions N and target quantiles.

q	0.1		0.3		0.499		0.7		0.9	
	H-FF	OF	H-FF	OF	H-FF	F	H-FF	OF	H-FF	OF
50	0,144	0,144	0,197	0,198	0,245	0,246	0,220	0,220	0,176	0,175
100	0,104	0,103	0,146	0,146	0,160	0,161	0,157	0,159	0,122	0,122
150	0,074	0,075	0,121	0,122	0,135	0,137	0,128	0,131	0,100	0,101
200	0,069	0,068	0,106	0,107	0,117	0,120	0,113	0,115	0,088	0,089
250	0,063	0,063	0,096	0,097	0,106	0,109	0,102	0,106	0,081	0,083
300	0,055	0,056	0,089	0,090	0,098	0,104	0,096	0,102	0,080	0,082
350	0,051	0,052	0,083	0,085	0,091	0,097	0,094	0,099	0,081	0,084
400	0,050	0,050	0,078	0,081	0,088	0,095	0,091	0,098	0,082	0,086
450	0,046	0,047	0,075	0,077	0,083	0,091	0,089	0,098	0,083	0,087
500	0,044	0,044	0,072	0,075	0,082	0,091	0,088	0,097	0,084	0,089

TABLE II: The estimation error for the OF and H-FF algorithms for the Normal distribution and for different values of the resolutions N and target quantiles.

q	0.1		0.3		0.499		0.7		0.9	
	H-FF	OF	H-FF	OF	H-FF	F	H-FF	OF	H-FF	OF
50	0,341	0,339	0,376	0,377	0,361	0,358	0,377	0,376	0,956	0,956
100	0,259	0,259	0,258	0,260	0,251	0,250	0,259	0,258	1,030	1,042
150	0,235	0,239	0,210	0,213	0,205	0,203	0,212	0,212	1,082	1,096
200	0,229	0,236	0,188	0,192	0,176	0,175	0,190	0,191	1,122	1,133
250	0,233	0,244	0,171	0,175	0,157	0,156	0,170	0,175	1,154	1,170
300	0,242	0,258	0,161	0,165	0,144	0,142	0,160	0,168	1,187	1,204
350	0,254	0,272	0,152	0,162	0,133	0,129	0,152	0,159	1,216	1,237
400	0,273	0,293	0,148	0,155	0,124	0,120	0,148	0,158	1,245	1,273
450	0,290	0,310	0,143	0,155	0,116	0,113	0,144	0,154	1,277	1,302
500	0,305	0,329	0,142	0,154	0,112	0,109	0,142	0,152	1,303	1,332

TABLE III: The estimation error for the OF and H-FF algorithms for the Exponential distribution and for different values of the resolutions N and target quantiles.

q	0.1		0.3		0.499		0.7		0.9	
	H-FF	OF	H-FF	OF	H-FF	F	H-FF	OF	H-FF	OF
50	0,159	0,158	0,253	0,254	0,335	0,332	0,399	0,401	0,473	0,464
100	0,109	0,109	0,181	0,182	0,235	0,237	0,285	0,290	0,378	0,385
150	0,078	0,078	0,149	0,148	0,193	0,198	0,237	0,247	0,370	0,381
200	0,074	0,073	0,129	0,130	0,169	0,174	0,215	0,227	0,386	0,404
250	0,066	0,066	0,116	0,117	0,153	0,160	0,204	0,219	0,416	0,442
300	0,057	0,058	0,107	0,109	0,141	0,152	0,200	0,218	0,459	0,489
350	0,056	0,056	0,099	0,102	0,134	0,147	0,195	0,219	0,501	0,540
400	0,053	0,053	0,095	0,097	0,130	0,144	0,197	0,223	0,544	0,587
450	0,048	0,048	0,090	0,094	0,125	0,142	0,199	0,228	0,598	0,639
500	0,047	0,048	0,088	0,091	0,122	0,142	0,203	0,237	0,638	0,687

TABLE IV: The estimation error for the OF and H-FF algorithms for the Chi-squared distribution and for different values of the resolutions N and target quantiles.

q	0.1		0.3		0.499		0.7		0.9	
	H-FF	OF	H-FF	OF	H-FF	F	H-FF	OF	H-FF	OF
50	0,088	0,088	0,254	0,254	0,348	0,345	0,453	0,454	0,600	0,606
100	0,063	0,063	0,149	0,149	0,234	0,231	0,322	0,326	0,519	0,525
150	0,051	0,052	0,126	0,125	0,192	0,192	0,270	0,272	0,535	0,567
200	0,045	0,045	0,105	0,104	0,167	0,170	0,245	0,253	0,597	0,638
250	0,040	0,040	0,094	0,095	0,150	0,153	0,227	0,243	0,686	0,731
300	0,037	0,036	0,085	0,085	0,139	0,142	0,220	0,238	0,765	0,822
350	0,033	0,033	0,079	0,079	0,129	0,136	0,218	0,239	0,842	0,915
400	0,031	0,031	0,074	0,075	0,122	0,128	0,220	0,244	0,933	0,987
450	0,029	0,029	0,070	0,070	0,118	0,125	0,218	0,254	1,003	1,062
500	0,027	0,027	0,067	0,068	0,113	0,121	0,222	0,258	1,073	1,134




Fig. 1: This figure depicts the variation of the estimation error for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n for the quantile of 30% and for $N = 150$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

it came to the convergence, the H-FF was faster than the OF in all settings. For instance, in Figure 1a, it took the H-FF around 20 iterations to reduce the error to a value under 0.02 while it took around 34 iterations for the OF to reach the same result – which is almost double the number of iterations.




Fig. 2: This figure depicts the variation of the estimation error for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n for the quantile of 46% and for $N = 150$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

B. Dynamic Environment

In this section, we report the simulation results for the case of dynamic environments for different values of the quantiles and for different values of the resolution parameter, N . In all the cases, we used the same values, namely $N = 400$, $N = 800$ and $N = 2,000$.

In order to model a dynamic environment, we modified the true quantile value after a fixed number of iterations. In the first set of experiments, the quantile




Fig. 3: This figure depicts the variation of the estimation error for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n for the quantile of 53% and for $N = 150$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

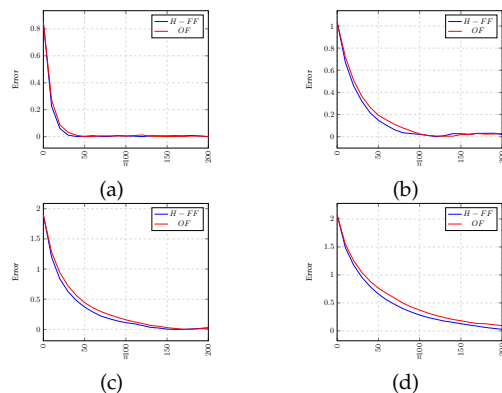


Fig. 4: This figure depicts the variation of the estimation error for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n for the quantile of 85% and for $N = 150$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

to be estimated cycled between the values 60%, 75%, 97%, 10%, and 80% after every 400th iterations. Figure 6 reports the evolution of the estimate for a resolution of $N = 400$ for the four distributions in question.

We also increased N to 800 and 2,000 in Figure 7 and Figure 8 while maintaining the same periodicity.

We observed that our scheme tracked the changing quantile for all the four distributions but that the rate of convergence after a change of the quantile value depended strongly on the choice of the resolution, N . We observed that for $N = 2,000$ in Figure 8, both estimators were not able to converge before the quantile changed its value, leading to a higher estimation error.

However, it is worth mentioning that even for a low resolution as low as $N = 400$, both estimators (OF and H-FF) encountered problems when there was a “large hop”. For example, in Figure 6a, we observed that between




Fig. 5: This figure depicts the variation of the estimation error for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n for the quantile of 97% and for $N = 150$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.




Fig. 6: This figure depicts the variations of the estimates for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n in a dynamically changing environment where the value changed every 400th time instant, and where the resolution parameter was $N = 400$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

time instants 1,200 and 1,600, the quantile changed significantly from 97% to 10% and thus, the H-FF was impaired just as in the case of the OF.

We also observed that the estimator converged the slowest for the Chi-Square distribution. The main reason for this is that the quantiles corresponding to the 60%, 75%, 97%, 10%, 80% quantiles for the Chi-square distribution are relatively distant from each other when compared to the other distributions. For the Chi-square distribution, the “dynamic” quantile values were 0.708, 1.323, 4.709, 0.0157 and 1.642 which are more distant from each other than, for example, the uniform distribution which in this case was 0.6, 0.75, 0.97, 0.1 and 0.8.

To display the results for other settings, we also doubled the pace by which the true quantile changed from




Fig. 7: This figure depicts the variations of the estimates for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n in a dynamically changing environment where the value changed every 400th time instant, and where the resolution parameter was $N = 800$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.




Fig. 8: This figure depicts the variations of the estimates for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n in a dynamically changing environment where the value changed every 400th time instant, and where the resolution parameter was $N = 2,000$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

every 400th iteration to every 200th iteration. Figure 9 depicts the dynamic of the estimates for the case when the change occurred at 400 and $N = 400$. We increased N while maintaining the same periodicity to 800 and 2,000 in Figure 7 and 8 respectively.

As expected both estimators faced more challenges in tracking than for the case where the periodicity was 400. However, the most important remark is that *in all the experiments* the H-FF outperformed the OF and was also faster to adapt to changes in the dynamic environments.




Fig. 9: This figure depicts the variations of the estimates for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n in a dynamically changing environment where the value changed every 200^{th} time instant, and where the resolution parameter was $N = 400$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.




Fig. 10: This figure depicts the variations of the estimates for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n in a dynamically changing environment where the value changed every 200^{th} time instant, and where the resolution parameter was $N = 800$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

IV. CONCLUSION

In this paper, we have dealt with the problem of estimating the quantiles of a distribution, which is a problem that is significantly more difficult than that of estimating the mean or a central/non-central moment. The use of these quantiles in pattern classification has become more widespread because they are more robust and applicable for non-parametric methods. The estimation of the quantiles is even more pertinent when one is mining (“infinite”) data streams, and is far more complex than the estimation of the moments because the increased complexity is more relevant as the size




Fig. 11: This figure depicts the variations of the estimates for the Original Frugal algorithm (OF) and our Higher-Fidelity Frugal (H-FF) with time n in a dynamically changing environment where the value changed every 200^{th} time instant, and where the resolution parameter was $N = 2,000$ for (a) the *Uniform* distribution, (b) the *Normal* distribution, (c) the *Exponential* distribution, and (d) the *Chi – Square* distribution.

of the data increases. We have focused on developing *incremental* quantile estimators [1], [2], which resort to updating the quantile estimates based on the most recent observation(s), leading to a very small computational and memory footprint.

This paper describes a scheme which is a confluence of three paradigms, namely, working with the foundations of Stochastic Point Location (SPL), the discretized world, and estimation of the quantiles in an incremental manner. We present a new quantile estimator which merges all these three concepts, and which we refer to as a Higher-Fidelity Frugal [1] (H-FF) quantile estimator. We have shown that the H-FF represents a substantial advancement of the family of Frugal estimators introduced in [1], and in particular to the so-called Original Frugal (OF) estimator.

Extensive simulation results show that our estimator outperforms the OF algorithm in terms of both speed and accuracy, and for both static and dynamic environments.

There are different extensions that can be envisaged for future work:

- In this paper, we worked within the domain of finite Markov chains, and have assumed that the true quantile was in the interval $[a, b]$. As a future work, we plan to extend the proof to infinite state Markov chains, which will, by no means, be trivial.
- The existing algorithm for quantile estimation was designed for data elements that were added one by one. A possible extension is to generalize our algorithm to handle not only data insertions, but also dynamic data operations such as deletions and updates, as proposed in [39].
- An interesting research direction is to attempt to

simultaneously estimate *multiple* quantile values. To achieve this, our present scheme will have to be modified so as to guarantee the monotonicity property of the quantiles, i.e., to maintain multiple quantile estimates while simultaneously ensuring that the estimates do not violate the monotonicity property. Although some preliminary results are currently available, they have yet to be perfected before they can be considered for publication.

REFERENCES

- [1] Q. Ma, S. Muthukrishnan, and M. Sandler, "Frugal streaming for estimating quantiles," in *Space-Efficient Data Structures, Streams, and Algorithms*. Springer, 2013, pp. 77–96.
- [2] A. Yazidi and H. Hammer, "Quantile estimation using the theory of stochastic learning," in *Proceedings of the 2015 Conference on research in adaptive and convergent systems*. ACM, 2015, pp. 7–14.
- [3] B. J. Oommen, "Stochastic searching on the line and its applications to parameter learning in nonlinear optimization," *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-27B, pp. 733–739, 1997.
- [4] B. J. Oommen and G. Raghunath, "Automata learning and intelligent tertiary searching for stochastic point location," *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-28B, pp. 947–954, 1998.
- [5] B. J. Oommen, G. Raghunath, and B. Kuipers, "Parameter learning from stochastic teachers and stochastic compulsive liars," *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-36B, pp. 820–836, 2006.
- [6] D.-S. Huang and W. Jiang, "A general cpl-ads methodology for fixing dynamic parameters in dual environments," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 42, no. 5, pp. 1489–1500, October 2012.
- [7] A. Yazidi, O. Granmo, B. John Oommen, and M. Goodwin, "A novel strategy for solving the stochastic point location problem using a hierarchical searching scheme," *IEEE Transactions on Cybernetics*, vol. 44, no. 11, pp. 2202–2220, Nov 2014.
- [8] M. A. L. Thathachar and B. J. Oommen, "Discretized reward-inaction learning automata," *Journal of Cybernetics and Information Science*, pp. 24–29, Spring 1979.
- [9] M. Agache and B. J. Oommen, "Generalized pursuit learning schemes: New families of continuous and discretized learning automata," *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*, vol. 32, no. 6, pp. 738–749, December 2002.
- [10] O.-C. Granmo, B. Oommen, S. Myrer, and M. Olsen, "Learning automata-based solutions to the nonlinear fractional knapsack problem with applications to optimal resource allocation," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 37, no. 1, pp. 166–175, feb. 2007.
- [11] O.-C. Granmo and B. J. Oommen, "Solving stochastic nonlinear resource allocation problems using a hierarchy of twofold resource allocation automata," *IEEE Transactions on Computers*, vol. 59, pp. 545–560, 2009.
- [12] J. K. Lanctôt and B. J. Oommen, "Discretized estimator learning automata," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-22, no. 6, pp. 1473–1483, November/December 1992.
- [13] B. J. Oommen and E. Hansen, "The asymptotic optimality of discretized linear reward-inaction learning automata," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-14, no. 3, May/June 1986.
- [14] B. J. Oommen, "Absorbing and ergodic discretized two-action learning automata," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-16, pp. 282–293, March/April 1986.
- [15] B. J. Oommen and M. Agache, "Continuous and discretized pursuit learning schemes: Various algorithms and their comparison," *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*, vol. 31, pp. 277–287, 2001.
- [16] B. J. Oommen and J. K. Lanctôt, "Discretized pursuit learning automata," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-20, no. 4, pp. 931–938, July/August 1990.
- [17] H. Robbins and S. Monro, "A stochastic approximation method," *The annals of mathematical statistics*, pp. 400–407, 1951.
- [18] H. Kushner and G. G. Yin, *Stochastic approximation and recursive algorithms and applications*. Springer Science & Business Media, 2003, vol. 35.
- [19] V. R. Joseph, "Efficient robbins–monro procedure for binary data," *Biometrika*, vol. 91, no. 2, pp. 461–470, 2004.
- [20] L. Tierney, "A space-efficient recursive procedure for estimating a quantile of an unknown distribution," *SIAM Journal on Scientific and Statistical Computing*, vol. 4, no. 4, pp. 706–711, 1983.
- [21] J. I. Munro and M. S. Paterson, "Selection and sorting with limited storage," *Theoretical computer science*, vol. 12, no. 3, pp. 315–323, 1980.
- [22] G. Cormode and S. Muthukrishnan, "An improved data stream summary: the count-min sketch and its applications," *Journal of Algorithms*, vol. 55, no. 1, pp. 58–75, 2005.
- [23] B. Weide, "Space-efficient on-line selection algorithms," in *Computer Science and Statistics: Proceedings of the Eleventh Annual Symposium on the Interface*, 1978, pp. 308–311.
- [24] A. Arasu and G. S. Manku, "Approximate counts and quantiles over sliding windows," in *Proceedings of the twenty-third ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, 2004, pp. 286–296.
- [25] M. Greenwald and S. Khanna, "Space-efficient online computation of quantile summaries," in *ACM SIGMOD Record*, vol. 30, no. 2. ACM, 2001, pp. 58–66.
- [26] S. Guha and A. McGregor, "Stream order and order statistics: Quantile estimation in random-order streams," *SIAM Journal on Computing*, vol. 38, no. 5, pp. 2044–2059, 2009.
- [27] F. Chen, D. Lambert, and J. C. Pinheiro, "Incremental quantile estimation for massive tracking," in *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2000, pp. 516–522.
- [28] A. Yazidi, O.-C. Granmo, and B. J. Oommen, "A stochastic search on the line-based solution to discretized estimation," in *Advanced Research in Applied Artificial Intelligence*. Springer, 2012, pp. 764–773.
- [29] A. Yazidi and B. J. Oommen, "Novel discretized weak estimators based on the principles of the stochastic search on the line problem," *IEEE transactions on cybernetics*, vol. 46, no. 12, pp. 2732–2744, 2016.
- [30] A. Yazidi, B. J. Oommen, G. Horn, and O.-C. Granmo, "Stochastic discretized learning-based weak estimation: a novel estimation method for non-stationary environments," *Pattern Recognition*, vol. 60, pp. 430–443, 2016.
- [31] B. J. Oommen, "Stochastic searching on the line and its applications to parameter learning in nonlinear optimization," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 27, no. 4, pp. 733–739, 1997.
- [32] J. M. Chambers, D. A. James, D. Lambert, and S. V. Wiel, "Monitoring networked applications with incremental quantile estimation," *Statistical Science*, pp. 463–475, 2006.
- [33] R. S. Boyer and J. S. Moore, *MJRTY—A Fast Majority Vote Algorithm*. Dordrecht: Springer Netherlands, 1991, pp. 105–117. [Online]. Available: https://doi.org/10.1007/978-94-011-3488-0_5
- [34] B. W. Schmeiser and S. J. Deutsch, "Quantile estimation from grouped data: The cell midpoint," *Communications in Statistics-Simulation and Computation*, vol. 6, no. 3, pp. 221–234, 1977.
- [35] O. Arandjelovic, D.-S. Pham, and S. Venkatesh, "Two maximum entropy-based algorithms for running quantile estimation in non-stationary data streams," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 25, no. 9, pp. 1469–1479, 2015.
- [36] R. Jain and I. Chlamtac, "The p 2 algorithm for dynamic calculation of quantiles and histograms without storing observations," *Communications of the ACM*, vol. 28, no. 10, pp. 1076–1085, 1985.
- [37] V. Naumov and O. Martikainen, "Exponentially weighted simultaneous estimation of several quantiles," *World Academy of Science, Engineering and Technology*, vol. 8, pp. 563–568, 2007.
- [38] J. Cao, L. E. Li, A. Chen, and T. Bu, "Incremental tracking of multiple quantiles for network monitoring in cellular networks," in *Proceedings of the 1st ACM workshop on Mobile internet through cellular networks*. ACM, 2009, pp. 7–12.
- [39] J. Cao, L. Li, A. Chen, and T. Bu, "Tracking quantiles of network data streams with dynamic operations," in *IEEE INFOCOM*. IEEE, 2010, pp. 1–5.