





Beyond Logarithmic Bounds: Querying in Constant Expected Time with Learned Indexes

Luis Alberto Croquevielle  

Imperial College London, UK

Guang Yang  

Imperial College London, UK

Liang Liang  

Imperial College London, UK
EPFL, Lausanne, Switzerland

Ali Hadian  

Imperial College London, UK

Thomas Heinis  

Imperial College London, UK

Abstract

Learned indexes leverage machine learning models to accelerate query answering in databases, showing impressive practical performance. However, theoretical understanding of these methods remains incomplete. Existing research suggests that learned indexes have superior asymptotic complexity compared to their non-learned counterparts, but these findings have been established under restrictive probabilistic assumptions. Specifically, for a sorted array with n elements, it has been shown that learned indexes can find a key in $O(\log(\log n))$ expected time using at most linear space, compared with $O(\log n)$ for non-learned methods.

In this work, we prove $O(1)$ expected time can be achieved with at most linear space, thereby establishing the tightest upper bound so far for the time complexity of an asymptotically optimal learned index. Notably, we use weaker probabilistic assumptions than prior research, meaning our work generalizes previous results. Furthermore, we introduce a new measure of statistical complexity for data. This metric exhibits an information-theoretical interpretation and can be estimated in practice. This characterization provides further theoretical understanding of learned indexes, by helping to explain why some datasets seem to be particularly challenging for these methods.

2012 ACM Subject Classification Theory of computation → Predecessor queries

Keywords and phrases Learned Indexes, Expected Time, Stochastic Processes, Rényi Entropy

Digital Object Identifier 10.4230/LIPIcs.ICDT.2025.19

Related Version *Extended Version*:

Querying in Constant Expected Time with Learned Indexes [6]

Funding This work is funded by DNAMIC (grant 101115389), NEO (grant 101115317), and ANID Chile through the Scholarship Program (DOCTORADO BECAS CHILE/2023 - 72230222).

1 Introduction

Query answering is of central importance in database systems, and developing efficient algorithms for this task is a key research problem. These algorithms often use an index structure to access records faster, at the cost of higher space usage to store the index. Arguably, the most common scenario in query answering involves point queries, where the aim is to retrieve all tuples where an attribute has an exact value [12, 21]. Also common are range queries, which ask for all tuples where an attribute is within a given interval.



© Luis Alberto Croquevielle, Guang Yang, Liang Liang, Ali Hadian, and Thomas Heinis; licensed under Creative Commons License CC-BY 4.0


28th International Conference on Database Theory (ICDT 2025).

Editors: Sudeepa Roy and Ahmet Kara; Article No. 19; pp. 19:1–19:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Predictive and corrective steps of a learned index.

Solutions based on hash tables can answer point queries efficiently [5], but are ineffective in handling range queries. To generalize well to these cases, the relevant attribute is usually stored as a sorted array A . In this configuration, answering point and range queries essentially comes down to finding an element in a sorted array. A point query for value q can be answered by finding the position of q in the array if it exists. A range query matching an interval $[q, q']$ can be answered by finding the first element bigger than or equal to q , and then scanning the array A from that position until an element greater than q' is found, or the array ends.

Since the array A is sorted, binary search can be used to find q or the first element greater than q . If A has n elements (also called *keys*), binary search takes $O(\log n)$ operations. This is asymptotically optimal in a random-access machine (RAM) model of computation, for comparison-based searching [17]. However, other index structures such as B+Trees [13, 19] are preferred in practice. They have the same asymptotic complexity but make use of hardware properties to improve performance.

It was suggested in [18] that classical index structures can be understood as models that take a key as input and predict its position in the sorted array. From this perspective, indexes such as B+Tree are potentially suboptimal because they do not take advantage of patterns in the data to improve performance. This has led to the introduction of learned indexes [18, 16, 8, 10, 11], which use machine learning models to predict the position of a key.

Learned indexes combine predictive and corrective steps, as illustrated in Figure 1. In the predictive step, a machine learning model is used to estimate the position of a key. Since this estimation may be wrong, a corrective step finds the exact position, usually by searching around the position estimated by the model. In principle, a learned index can improve upon classical methods by limiting the search to a small range, determined by the prediction error, rather than having to search the entire array. Most learned indexes use simple machine learning models, such as piecewise linear functions, which are fast to compute.

While learned indexes exhibit impressive experimental performance [8, 11, 14, 16], their theoretical understanding is still at an early stage. Existing research shows that specific learned indexes can achieve better asymptotic space or time complexity than traditional structures [34, 9], but these results rely on disparate and potentially restrictive assumptions. Furthermore, some datasets are notoriously challenging for learned indexes [24], and no theoretical analysis has yet offered an adequate explanation for this.

In this work, we analyze the asymptotic expected complexity of learned indexes. Specifically, we prove that learned indexes can achieve constant expected query time with at most linear space, constituting the best upper bound so far for the performance of an asymptotically optimal learned index. Our proof is constructive, meaning we provide a specific index that achieves this complexity. This type of index has been considered before [34], but we

contribute tighter complexity bounds than previously available. Compared to prior research, our complexity bounds hold under a more general probabilistic model, and in particular, our results do not require any independence assumptions for the data generating process.

We also address the question of why certain datasets are challenging for learned indexes. It has been suggested that, in those cases, the data distribution is inherently hard to learn [24], and several metrics have been proposed to capture this difficulty [34, 9, 32]. We introduce a new measure of data complexity ρ_f with better mathematical properties. Our metric is related to Rényi entropy in information theory, helping to provide a theoretical grounding for ρ_f . Also, ρ_f can be statistically estimated, making it potentially useful in practice to predict *a priori* how well a learned index will perform. Our experiments support this idea.

2 Preliminaries

2.1 Data

Learned indexes use statistical models. Therefore, theoretical analysis of learned indexes is done under a probabilistic model of some kind [34, 9]. In general, keys are assumed to come from a data generating process with specified properties. For instance, in [34] it is assumed that the keys are independent and identically distributed, while [9] uses a completely different model. We strive to generalize the probabilistic setting as much as possible.

Formally, we model a database attribute X as a stochastic process $X = \{X_i\}_{i \in \mathbb{N}}$. We assume all random variables X_i have the same cumulative distribution function F . For any time $n \in \mathbb{N}$, we can observe the data X_1, X_2, \dots, X_n generated so far. Sorting the $\{X_i\}_{i=1}^n$ gives rise to a new set of random variables, usually called the *order statistics* [7], which we denote $X_{(1)}, X_{(2)}, \dots, X_{(n)}$. We consider the sorted array A at time n to be formed by this last group of random variables, such that $A[i] = X_{(i)}$ for all $i = 1, \dots, n$. Duplicate keys are possible in principle, and do not constitute a problem. In that case, $X_{(i+1)} = X_{(i)}$ for some i and all analysis remains the same.

That way, in our setting a stochastic process generates the data in attribute X . For each $n \in \mathbb{N}$, there is a sorted array A with the first n keys from this process, and a learned index can be built to search this array. For now, we do not assume anything about F . Also, we do not assume independence of the $\{X_i\}_{i=1}^n$ variables. In Sections 3 and 4, we analyze our results under certain probabilistic assumptions. In Section 5, we compare our model and results to those of prior work.

2.2 Learning Problem

Assume array A has n keys. The learning problem can be formalized by the introduction of a **rank** function. For any number $q \in \mathbb{R}$, $\mathbf{rank}_A(q)$ is defined as the number of elements in A that are smaller than or equal to q , that is

$$\mathbf{rank}_A(q) = \sum_{i=1}^n \mathbf{1}_{X_{(i)} \leq q},$$

where $\mathbf{1}$ is the indicator function. We omit the subscript A if it is clear from context. For a point query with value q , $\mathbf{rank}(q)$ gives the position of q in the array if it exists. For a range query matching an interval $[q, q']$, the relevant records are found by scanning the array A , starting at position $\mathbf{rank}(q)$ and ending when an element greater than q' is found. Since **rank** is sufficient to answer these queries, we focus on the problem of learning the **rank** function under specific probabilistic assumptions.

2.3 Model of Computation

The model of computation we use is the Random Access Machine (RAM) under the uniform cost criterion [1]. In this model, each instruction requires one unit of time and each register (storing an arbitrary integer) uses one unit of space. This is in contrast to part of the existing literature. In [34] the computation model is such that $O(\log n)$ space units are needed to store an integer n . On the other hand, in [9] an I/O model of computation is used [20], which assumes the existence of an external memory (e.g., disk) from which data is read and written to. Under this model, time complexity consists only of the number of I/O operations.

We argue the RAM model under the uniform cost criterion is better suited for analysing learned indexes than either of these two. On the one hand, numbers are usually represented as data types with a fixed number of bits (e.g., `long int` in C++). On the other hand, most learned indexes work in main memory and avoid the use of disk storage. To make results comparable, in Section 5 we restate all previous complexity results using our RAM model.

Regarding space complexity, we take the standard approach of only counting the redundant space introduced by the index, and not the space necessary to store the array A in the first place [10]. Regarding time complexity, we do not consider the time needed to train the model, just for finding the keys. We focus the asymptotic analysis on *expected* time, where the expectation is taken over the realizations of the stochastic process and the choice of query.

For any $n \in \mathbb{N}$, let R_n be a procedure that takes the sorted array $A = [X_{(1)}, \dots, X_{(n)}]$ as input, and outputs an index structure $R_n(A)$. Denote by $S(R_n(A))$ the space overhead of $R_n(A)$. Also, for any $q \in \mathbb{R}$ denote by $T(R_n(A), q)$ the query time for q over array A when using the index $R_n(A)$. We are interested in bounds for S and T that do not depend on A or q . With that in mind, we denote

$$\bar{S}_n = \sup_{A=[X_{(1)}, \dots, X_{(n)}]} S(R_n(A)), \quad \bar{T}_n = \mathbb{E}[T(R_n(A), q)],$$

where \mathbb{E} denotes expected value. In Section 3.1 we state asymptotic results for \bar{S}_n and \bar{T}_n .

3 Main Results and ESPC Index

3.1 Complexity Bounds

We mentioned above that all the X_i share the same cumulative distribution function. Formally, we consider that the X_i are defined in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, such that $F_{X_i}(x) := \mathbb{P}(X_i \leq x)$ is the same for all i , and is denoted by F . Now, assume F can be characterized by a square-integrable density function f and define ρ_f as:

$$\rho_f = \|f\|_2^2 = \int_{\mathbb{R}} f^2(x) dx.$$

Moreover, assume f has bounded support $[a, b]$, that is, $f(x) = 0$ for all $x \notin [a, b]$. This can be a reasonable assumption in many practical applications [22, 23] where there are known bounds to the values of data: for instance, for storing medical information such as test results, or personal financial information. Our results can be extended to the case of f with unbounded support, by incurring a penalty on space usage. In Section 4.3.3 we show how this penalty relates to the tails of the probability distribution.

The main complexity results are stated in Theorems 1 and 2. Both are attained by the ESPC index, which we introduce in Section 3.2. Both theorems as stated here assume that the query parameter follows the same probability distribution as the keys. This is done

for ease of presentation. In Section 4.3.2 we prove that these results can be generalized to include the case when that assumption does not hold.

► **Theorem 1.** *Suppose f has support $[a, b]$ and $\rho_f < \infty$. Define $\rho = \log((b - a)\rho_f)$. Then, there is a procedure R_n for building learned indexes such that $\overline{S}_n = O(n)$ and $\overline{T}_n = O(\rho)$. That is, for array A with n keys an index can be built with space overhead $O(n)$ and expected query time $O(\rho)$. Since ρ is independent of n , expected time is asymptotically $O(1)$.*

► **Theorem 2.** *Suppose f has support $[a, b]$ and $\rho_f < \infty$. Define $\rho = \log((b - a)\rho_f)$. Then, there is a procedure R_n for building learned indexes such that $\overline{S}_n = O(\frac{n}{\log n})$ and $\overline{T}_n = O(\rho + \log(\log n))$. That is, for array A with n keys an index can be built with space overhead $O(\frac{n}{\log n})$ and expected query time $O(\rho + \log(\log n))$. Since ρ is independent of n , expected time is asymptotically $O(\log(\log n))$.*

We prove Theorems 1 and 2 in Section 4.3.1, as corollaries of a more general result.

3.2 Equal-Split Piecewise Constant Index

All complexity results in Section 3.1 are achieved by the *Equal-Split Piecewise Constant* (ESPC) index, which we now describe. As explained in Section 1, learned indexes usually combine predictive and corrective steps. The ESPC index follows this idea, and is remarkably simple to define and implement. It is important to mention that a similar design, the PCA Index, was analyzed in [34], but the complexity bounds we prove for this type of index are tighter and more general. Moreover, there are some design differences. First, the ESPC index uses exponential search for the corrective step, while the PCA Index uses binary search. More importantly, the PCA Index presumes knowledge of the interval $[a, b]$ where the distribution is supported, a potentially restrictive assumption that we dispense with for the ESPC index.


The general idea is to learn a function $\hat{\mathbf{r}}$ to approximate the **rank** function up to some error. Then **rank** can be computed exactly by first evaluating $\hat{\mathbf{r}}$ and then correcting the error by use of an exponential search algorithm [2]. If the prediction error is small, this gives a procedure for computing **rank** exactly in low expected time. The approximator function $\hat{\mathbf{r}}$ is built as a piecewise constant function defined over equal-length subintervals. This equal-length property ensures that $\hat{\mathbf{r}}$ can be evaluated in constant time.

We now describe how the approximator $\hat{\mathbf{r}}$ is defined. Take an array A formed by the sorted keys $X_{(1)}, \dots, X_{(n)}$. For a positive integer K , divide the range $[X_{(1)}, X_{(n)}]$ into K subintervals of equal length, which we denote by $\{I_k\}_{k=1}^K$. Figure 2 illustrates this type of partition for a dataset with 40 keys and $K = 4$. Formally, for each $k = 1, \dots, K$ the k -th subinterval is defined as $I_k = [t_{k-1}, t_k]$ where

$$t_k = X_{(1)} + k\delta \quad \text{with} \quad \delta = \frac{X_{(n)} - X_{(1)}}{K}.$$

Notice that all intervals have the same length, but not necessarily the same number of keys. Now, we want $\hat{\mathbf{r}}$ to be defined as a piecewise constant function. If $q \in \mathbb{R}$ does not belong to any I_k , this means that either $q < X_{(1)}$ or $q > X_{(n)}$, so that $\mathbf{rank}(q) = 0$ or $\mathbf{rank}(q) = n$, respectively. On the other hand, for each subinterval I_k the ESPC index stores a constant

19:6 Querying in Constant Expected Time with Learned Indexes



■ **Figure 2** Partition of key range into four equal-length intervals, with approximator of **rank**.

value \hat{r}_k which approximates **rank** over I_k . Combining these ideas, we define $\hat{\mathbf{f}}$ as

$$\hat{\mathbf{f}}(q) = \begin{cases} 0 & \text{if } q < X_{(1)} \\ \hat{r}_1 & \text{if } q \in I_1 \\ \vdots & \\ \hat{r}_K & \text{if } q \in I_K \\ n & \text{if } q > X_{(n)}. \end{cases}$$

Figure 2 represents this type of piecewise approximator over equal-length subintervals. It remains to define the $\{\hat{r}_k\}_{k=1}^K$ values. Since **rank** is an increasing function, for all $q \in I_k$ it holds $\mathbf{rank}(t_{k-1}) \leq \mathbf{rank}(q) \leq \mathbf{rank}(t_k)$. This motivates a candidate for \hat{r}_k in the form of

$$\hat{r}_k = \frac{1}{2} (\mathbf{rank}(t_{k-1}) + \mathbf{rank}(t_k)) = \mathbf{rank}(t_{k-1}) + \frac{n_k}{2},$$

where $n_k = \mathbf{rank}(t_k) - \mathbf{rank}(t_{k-1})$ is the number of keys in I_k . As we prove in Section 4.1, this means that \hat{r}_k can approximate **rank**(q) for any $q \in I_k$ with error at most $n_k/2$, so that the quality of $\hat{\mathbf{f}}$ as an approximator depends on the expected number of keys per interval. Algorithm 1 details the index construction, based on our definitions for $\{I_k\}$ and $\{\hat{r}_k\}$.

■ **Algorithm 1** Construction of ESPC index.

Input: Sorted array A , positive integer K
Output: ESPC index for A with K equal-length subintervals

1. $n \leftarrow \text{length}(A)$
2. $\delta \leftarrow (X_{(n)} - X_{(1)})/K$
3. $N \leftarrow$ array of length K with all zeros
4. // Compute number of keys in each subinterval, store in N
5. **for** $i = 1$ **to** n **do**
6. $k \leftarrow \lceil \frac{1}{\delta}(A[i] - X_{(1)}) \rceil$
7. $N[k] \leftarrow N[k] + 1$
8. **end**
9. $r \leftarrow$ array of length K
10. $r[1] \leftarrow \frac{1}{2}N[1]$
11. // Compute rank estimator \hat{r}_k for each subinterval, store in r
12. **for** $k = 2$ **to** K **do** $r[k] \leftarrow r[k-1] + \frac{1}{2}N[k]$;
13. **return** $R = (r, \delta)$

The approximator $\hat{\mathbf{r}}$ can be evaluated in constant time for any $q \in \mathbb{R}$. If $q < X_{(1)}$ then $\hat{\mathbf{r}}(q) = 0$, and if $q > X_{(n)}$ then $\hat{\mathbf{r}}(q) = n$. If neither is true, then q belongs to some interval I_k . In that case, $\hat{\mathbf{r}}(q) = \hat{r}_k$ where k is given by $k = \left\lceil K \frac{q - X_{(1)}}{X_{(n)} - X_{(1)}} \right\rceil = \left\lceil \frac{1}{\delta} (q - X_{(1)}) \right\rceil$. To find $\mathbf{rank}(q)$ exactly, the ESPC index then performs an exponential search on array A around position $\lceil \hat{\mathbf{r}}(q) \rceil$. We prove in Section 4 that the expected value of the approximation error $|\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|$ is low enough to establish the results of Section 3.1. Algorithm 2 formalizes this evaluation process. It assumes the existence of a procedure $\text{EXPONENTIALSEARCH}(A, i, q)$ which uses exponential search around index i to find the position of q in array A .

■ **Algorithm 2** Evaluation of ESPC index.

Input: Sorted array A , $q \in \mathbb{R}$, ESPC index $R = (r, \delta)$

Output: Exact value of $\mathbf{rank}(q)$

1. **if** $q < X_{(1)}$ **then return** 0;
 2. **else if** $q > X_{(n)}$ **then return** n ;
 3. **else**
 4. $k \leftarrow \left\lceil \frac{1}{\delta} (q - X_{(1)}) \right\rceil$
 5. $i \leftarrow \lceil r[k] \rceil$
 6. **return** $\text{EXPONENTIALSEARCH}(A, i, q)$
- end**
-

4 Complexity Proofs and Analysis

4.1 Preliminary Results


We now turn to an analysis of the space and time complexity of the ESPC index. In the following, assume the sorted array A has n elements and the index is built with K subintervals $\{I_k\}_{k=1}^K$. As seen from Algorithm 1, an ESPC index can be represented as a tuple $R = (r, \delta)$ where r is an array of length K storing the approximators $\{\hat{r}_k\}_{k=1}^K$ and δ is the length of the subintervals. For $q \in \mathbb{R}$ recall that $\hat{\mathbf{r}}(q)$ is the approximation for $\mathbf{rank}(q)$ before the exponential search (see Algorithm 2).

In the RAM model with the uniform cost criterion, each stored value uses 1 unit of space. Hence, the space used to store an ESPC index $R = (r, \delta)$ is $O(K)$. Additionally, as seen from Algorithm 2, the evaluation for any $q \in \mathbb{R}$ takes constant time, plus the exponential search. An exponential search takes time $O(2 \log \varepsilon) = O(\log \varepsilon)$ [2], where ε is the difference between the starting position and the correct position. This is summarized in the following result.

► **Proposition 3.** *Let A be a sorted array with n entries, and R be an ESPC index for A with K subintervals. Then, R uses $O(K)$ space and given $q \in \mathbb{R}$, the ESPC index R can be used to find $\mathbf{rank}(q)$ in $O(\log \varepsilon)$ time, where $\varepsilon = |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|$.*

In light of Proposition 3, a way to estimate the time complexity of the ESPC index is to bound the expected value of ε as a function of K . Hence, our focus is on finding an expression for the approximation error $\varepsilon = |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|$. We start with the following result, which was mentioned in passing in Section 3.2.

► **Lemma 4.** *Let $k \in \{1, \dots, K\}$ and $q \in I_k$. Then, it holds that $|\mathbf{rank}(q) - \hat{\mathbf{r}}(q)| \leq n_k/2$ where n_k is the number of keys that fall on interval I_k .*



■ **Figure 3** Equal-length partitions for $[X_{(1)}, X_{(n)}]$ (given by $\{I_k\}$) and $[a, b]$ (given by $\{J_k\}$).

The proof is straightforward and can be found in the extended version of the paper [6]. We now introduce other necessary notation. First, notice that if f has bounded support $[a, b]$, then $[X_{(1)}, X_{(n)}] \subseteq [a, b]$ holds with probability 1. Now, for a given K we can define a set of subintervals $\{J_k\}_{k=1}^K$ that partition the interval $[a, b]$ into K equal-length parts, similarly to how the $\{I_k\}_{k=1}^K$ form a partition of $[X_{(1)}, X_{(n)}]$. See Figure 3 for an illustration with $K = 3$.

From this definition, it may be the case that some I_k is not entirely contained within just one J_k interval, as Figure 3 illustrates for the case of I_2 . Notice however that if this happens, then I_k must be contained in the union of two consecutive J_k intervals, as exemplified by $I_2 \subseteq J_1 \cup J_2$ in Figure 3. We formalize this observation below, which is easily verified.

► **Fact 5.** *Suppose $[X_{(1)}, X_{(n)}] \subseteq [a, b]$. Take $k, \ell \in \{1, \dots, K\}$ such that $I_k \cap J_\ell \neq \emptyset$. Then either $I_k \subseteq J_{\ell-1} \cup J_\ell$ or $I_k \subseteq J_\ell \cup J_{\ell+1}$ holds. As a consequence, it also holds that $I_k \subseteq J_{\ell-1} \cup J_\ell \cup J_{\ell+1}$, where J_0, J_{K+1} are defined as the empty set \emptyset for completion.*

4.2 Main Probabilistic Bound

Recapping, by Proposition 3 the ESPC index time complexity for input q depends on the approximation error $\varepsilon = |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|$. By Lemma 4, the error ε critically depends on the number of keys in the subintervals $\{I_k\}$. We now state and prove the key results which provide bounds for the approximation error. As mentioned in Section 3.1, we assume that the density f has bounded support $[a, b]$ for some $a, b \in \mathbb{R}$. That is, $f(x) = 0$ for any $x \notin [a, b]$. Our results can be extended to the case of f with unbounded support by incurring a penalty in terms of space complexity. Notice that the $\{J_k\}$ implicitly define a random variable, which models how the $\{X_i\}_{i=1}^n$ are distributed among the subintervals. This can be characterized by a parameter (p_1, \dots, p_K) where $p_k = \mathbb{P}(X_i \in J_k)$. Notice that p_k is not indexed by i because all X_i follow the same distribution. In this context, the core of the time complexity analysis is contained in the following theorem.

► **Theorem 6.** *Suppose f has support $[a, b]$. Let A be an array consisting of the $\{X_i\}_{i=1}^n$ sorted in ascending order. Let $\hat{\mathbf{r}}$ be the approximator of the ESPC index with K subintervals, as described in Section 3.2, and denote the approximation error for $q \in \mathbb{R}$ as $\varepsilon(q) = |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|$. Then, if the query parameter q has the same density function f as the $\{X_i\}$, it holds that $\mathbb{E}[\varepsilon] \leq \frac{3n}{2} \sum_{k=1}^K p_k^2$.*

Proof. We know that $\mathbb{E}[\varepsilon] = \mathbb{E}[\mathbb{E}[\varepsilon | X_1, \dots, X_n]]$, by the tower property for conditional expectation. The random nature of the error depends on both the choice of q and the values of $\{X_i\}_{i=1}^n$. To see this, recall that the variables X_1, \dots, X_n determine the value of both \mathbf{rank} and $\hat{\mathbf{r}}$. The inner expected value gives

$$\mathbb{E}[\varepsilon | X_1, \dots, X_n] = \int_{\mathbb{R}} \varepsilon(q) f(q) dq = \int_{X_{(1)}}^{X_{(n)}} \varepsilon(q) f(q) dq,$$

where the last equality comes from the fact that $\varepsilon(q) = 0$ if $q \notin [X_{(1)}, X_{(n)}]$. Replacing $\varepsilon(q) = |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|$ and noticing that the $\{I_k\}$ form a partition of $[X_{(1)}, X_{(n)}]$:

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] = \sum_{k=1}^K \int_{I_k} |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)| f(q) dq \leq \frac{1}{2} \sum_{k=1}^K n_k \int_{I_k} f(q) dq,$$

where n_k is the number of keys in subinterval I_k and we have used Lemma 4 to bound the error within each I_k . Now, denote by $n_{k,\ell}$ the number of keys in $I_k \cap J_\ell$. Then, we can write

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] \leq \frac{1}{2} \sum_{k=1}^K \left(\sum_{\ell=1}^K n_{k,\ell} \right) \int_{I_k} f(q) dq = \frac{1}{2} \sum_{\ell=1}^K \sum_{k=1}^K n_{k,\ell} \int_{I_k} f(q) dq, \quad (1)$$

where first we have written n_k as $\sum_{\ell=1}^K n_{k,\ell}$ and then we have exchanged the order of summation. Now, denote by $L(k)$ the set of ℓ indices such that $I_k \cap J_\ell \neq \emptyset$. Notice that $n_{k,\ell} = 0$ if $\ell \notin L(k)$. Hence, inequality (1) becomes:

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] \leq \frac{1}{2} \sum_{\ell=1}^K \sum_{k:\ell \in L(k)} n_{k,\ell} \int_{I_k} f(q) dq. \quad (2)$$

By Fact 5 we know that $I_k \subseteq J_{\ell-1} \cup J_\ell \cup J_{\ell+1}$ for all ℓ and k such that $\ell \in L(k)$. Using the notation $p_\ell = \mathbb{P}(X_i \in J_\ell) = \int_{J_\ell} f(q) dq$ introduced above, from inequality (2) we get

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] \leq \frac{1}{2} \sum_{\ell=1}^K (p_{\ell-1} + p_\ell + p_{\ell+1}) \sum_{k:\ell \in L(k)} n_{k,\ell} = \frac{1}{2} \sum_{\ell=1}^K (p_{\ell-1} + p_\ell + p_{\ell+1}) m_\ell, \quad (3)$$

where m_ℓ denotes the number of keys in J_ℓ . Now, applying expected value on both sides of inequality (3), we get that $\mathbb{E}[\varepsilon] = \mathbb{E}[\mathbb{E}[\varepsilon \mid X_1, \dots, X_n]] \leq \frac{1}{2} \sum_{\ell=1}^K (p_{\ell-1} + p_\ell + p_{\ell+1}) \mathbb{E}[m_\ell]$. Furthermore, we have

$$\mathbb{E}[m_\ell] = \mathbb{E} \left[\sum_{i=1}^n \mathbf{1}_{X_i \in J_\ell} \right] = \sum_{i=1}^n \mathbb{E}[\mathbf{1}_{X_i \in J_\ell}] = \sum_{i=1}^n \mathbb{P}(X_i \in J_\ell) = np_\ell.$$

Replacing in the expression for $\mathbb{E}[\varepsilon]$, we get the following inequality:

$$\mathbb{E}[\varepsilon] \leq \frac{1}{2} \sum_{\ell=1}^K (p_{\ell-1} + p_\ell + p_{\ell+1}) \mathbb{E}[m_\ell] = \frac{n}{2} \sum_{\ell=1}^K (p_{\ell-1} p_\ell + p_\ell^2 + p_\ell p_{\ell+1}). \quad (4)$$

Now, using Cauchy–Schwarz inequality and the fact that $p_0 = \mathbb{P}(X_i \in J_0) = 0$ we get:

$$\sum_{\ell=1}^K p_{\ell-1} p_\ell \leq \left(\sum_{\ell=1}^K p_{\ell-1}^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2} = \left(\sum_{\ell=1}^{K-1} p_\ell^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2} \leq \left(\sum_{\ell=1}^K p_\ell^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2} = \sum_{\ell=1}^K p_\ell^2.$$

Similarly, using Cauchy–Schwarz inequality and the fact that $p_{K+1}^2 = 0$ we can show that $\sum_{\ell=1}^K p_\ell p_{\ell+1} \leq \sum_{\ell=1}^K p_\ell^2$. Finally, applying these inequalities in (4), we get the desired result:

$$\mathbb{E}[\varepsilon] \leq \frac{n}{2} \sum_{\ell=1}^K (p_{\ell-1} p_\ell + p_\ell^2 + p_\ell p_{\ell+1}) \leq \frac{n}{2} \left(\sum_{\ell=1}^K p_\ell^2 + \sum_{\ell=1}^K p_\ell^2 + \sum_{\ell=1}^K p_\ell^2 \right) \leq \frac{3n}{2} \sum_{\ell=1}^K p_\ell^2. \quad \blacktriangleleft$$

Theorem 6 is the key tool for proving our results, as we show in Section 4.3.

4.3 Proof of Asymptotic Complexity

We now use Theorem 6 to prove the results stated in Section 3.1. For this, we need to relate the error bound in Theorem 6 with $\rho_f = \|f\|_2^2$. For this, notice that

$$p_k^2 = \left(\int_{J_k} f(x) dx \right)^2 \leq \left(\int_{J_k} 1 dx \right) \left(\int_{J_k} f^2(x) dx \right) = |J_k| \int_{J_k} f^2(x) dx,$$

where $|J_k|$ denotes the length of interval J_k and we have used Cauchy-Schwarz inequality to bound $(\int_{J_k} f(x) dx)^2$. Substituting this in the error bound from Theorem 6 and using the fact that $|J_k| = (b-a)/K$ for all k , we obtain

$$\mathbb{E}[\varepsilon] \leq \frac{3n}{2} \sum_{k=1}^K p_k^2 \leq \frac{3n}{2} \sum_{k=1}^K \frac{(b-a)}{K} \int_{J_k} f^2(x) dx = \frac{3(b-a)}{2} \frac{n}{K} \sum_{k=1}^K \int_{J_k} f^2(x) dx$$

Since $\{J_k\}$ is a partition of $[a, b]$ and f is square-integrable:

$$\mathbb{E}[\varepsilon] \leq \frac{3(b-a)}{2} \frac{n}{K} \int_a^b f^2(x) dx = \frac{3(b-a)}{2} \frac{n}{K} \int_{\mathbb{R}} f^2(x) dx = \frac{3(b-a)}{2} \rho_f \frac{n}{K}.$$

We can summarize the main error bound for the ESPC index in the following proposition.

► **Proposition 7.** *Under the conditions of Theorem 6, and assuming $\rho_f < \infty$, it holds*

$$\mathbb{E}[\varepsilon] \leq \frac{3(b-a)}{2} \rho_f \frac{n}{K}.$$

The complexity results from Section 3.1 are proved in Section 4.3.1 as corollaries of Proposition 7. In Sections 4.3.2 and 4.3.3 we discuss two important generalizations of our results. In Section 5, we present further analysis and a comparison to previous results.

4.3.1 Proofs of Main Complexity Bounds

Proof of Theorem 1. Let $n \in \mathbb{N}$ and consider the following procedure R_n for generating an index: given a sorted array $A = [X_{(1)}, \dots, X_{(n)}]$, the output $R_n(A)$ consists of an ESPC index built with $K = n$ subintervals. By Proposition 3, this means that space overhead is $S(R_n(A)) = O(K) = O(n)$. As this is independent of the specific values stored in A , we have $\bar{S}_n = O(n)$. On the other hand, by Theorem 6 and Proposition 7 we can bound the expected prediction error by

$$\mathbb{E}[\varepsilon] \leq \frac{3(b-a)}{2} \rho_f \frac{n}{n} = \frac{3(b-a)}{2} \rho_f.$$

Furthermore, Jensen's inequality implies

$$\mathbb{E}[\log \varepsilon] \leq \log \mathbb{E}[\varepsilon] \leq \log((b-a)\rho_f) + \log(3/2) \leq \log((b-a)\rho_f) + 1.$$

By Proposition 3, we know query time for q for the ESPC index is $O(\log \varepsilon(q))$. Accordingly, the expected query time is $\bar{T}_n = O(\mathbb{E}[\log \varepsilon]) = O(\log((b-a)\rho_f) + 1) = O(\log((b-a)\rho_f))$. This concludes the proof. ◀

Proof of Theorem 2. We use the same argument as the proof for Theorem 1, but setting $K = \frac{n}{\log n}$ subintervals instead of $K = n$. ◀

4.3.2 Distribution of query parameter

As can be seen from the statement of Theorem 6, we have assumed that q has the same probability density function f as the $\{X_i\}$. We believe this to be a reasonable approximation in many contexts, but our analysis can be extended to q otherwise distributed, which constitutes an important generalization with respect to previous work [34].

► **Proposition 8.** *Assume the conditions of Theorem 6. Additionally, suppose $\rho_f < \infty$ and that the search parameter q distributes according to a density function g . Further suppose that g is square-integrable and denote $\rho_g = \|g\|_2^2$. Then, it holds that*

$$\mathbb{E}[\varepsilon] \leq \frac{3(b-a)}{2} \sqrt{\rho_f \rho_g} \frac{n}{K}.$$

See Appendix A for the proof. Hence, for the general case where q does not distribute according to f , our results retain essentially the same form. For ease of presentation, we assume $g = f$ in subsequent analyses. By Proposition 8, this assumption entails no loss of generality; for q otherwise distributed, we substitute $\sqrt{\rho_f \rho_g}$ in place of ρ_f .

4.3.3 Unbounded support

Proposition 7 can be extended to f with unbounded support. For $a, b \in \mathbb{R}$ denote by $A_{a,b}$ the event that $[X_{(1)}, X_{(n)}] \not\subseteq [a, b]$. Take values of a, b such that $\mathbb{P}(A_{a,b}) \leq \frac{1}{\log n}$. Then, the expected query time is $O(\mathbb{E}[\log \varepsilon])$ where the expected value can be written as

$$\mathbb{E}[\log \varepsilon] = E_1 + E_2 \text{ where } E_1 = \mathbb{E}[\log \varepsilon \mid A_{a,b}^C] \mathbb{P}(A_{a,b}^C) \text{ and } E_2 = \mathbb{E}[\log \varepsilon \mid A_{a,b}] \mathbb{P}(A_{a,b}).$$

Since the prediction error ε can never exceed n , the E_2 term can be bounded as

$$E_2 = \mathbb{E}[\log \varepsilon \mid A_{a,b}] \mathbb{P}(A_{a,b}) \leq \log n \cdot \mathbb{P}(A_{a,b}) \leq \log n \cdot \frac{1}{\log n} = 1.$$

On the other hand, the event $A_{a,b}^C$ means that $[X_{(1)}, X_{(n)}] \subseteq [a, b]$. This allows us to apply Jensen's inequality along with Proposition 7 to get

$$E_1 = \mathbb{E}[\log \varepsilon \mid A_{a,b}^C] \mathbb{P}(A_{a,b}^C) \leq \mathbb{E}[\log \varepsilon \mid A_{a,b}^C] \leq \log \mathbb{E}[\varepsilon \mid A_{a,b}^C] \leq \log \left(\frac{3(b-a)}{2} \rho_f \frac{n}{K} \right).$$

Putting everything together, this means that expected query time is

$$O(\mathbb{E}[\log \varepsilon]) = O \left(\log \left(\frac{3(b-a)}{2} \rho_f \frac{n}{K} \right) + 1 \right) = O \left(\log \left((b-a) \rho_f \frac{n}{K} \right) \right). \quad (5)$$

Hence, time complexity depends on the values of a, b that guarantee $\mathbb{P}(A_{a,b}) \leq \frac{1}{\log n}$. These values depend on how heavy-tailed the distribution is. We propose the following two variations of Theorem 1, with other results being available based on different assumptions.

► **Theorem 9.** *Suppose the $\{X_i\}$ have finite mean μ and variance σ^2 , and that $\rho_f < \infty$. Then, there is a procedure R_n for building learned indexes with $\bar{S}_n = O(n\sqrt{n \log n})$ and $\bar{T}_n = O(\rho)$, where $\rho = \log(2\sigma\rho_f)$. Since ρ is independent of n , expected time is asymptotically $O(1)$.*

► **Theorem 10.** *Suppose $\rho_f < \infty$ and that the distribution of the $\{X_i\}$ is subexponential [30], meaning there is a constant $C > 0$ such that $\mathbb{P}(|X_i| \geq x) \leq 2e^{-Cx}$ for all $x \geq 0$. Then, there is a procedure R_n for building learned indexes with $\bar{S}_n = O(n \log n)$ and $\bar{T}_n = O(\rho)$, where $\rho = \log(4\rho_f/C)$. Since ρ is independent of n , expected time is asymptotically $O(1)$.*

19:12 Querying in Constant Expected Time with Learned Indexes

We include below the proof of Theorem 9. For Theorem 10, the proof is available in the extended version of the paper [6]. It can also be viewed as a corollary of Lemma 3.4 in [34].

Proof of Theorem 9. Set $a = \mu - \sigma\sqrt{n \log n}$ and $b = \mu + \sigma\sqrt{n \log n}$. It holds that

$$\mathbb{P}(A_{a,b}) = \mathbb{P}([X_{(1)}, X_{(n)}] \not\subseteq [a, b]) = \mathbb{P}\left(\bigcup_{i=1}^n \{X_i \notin [a, b]\}\right) \leq \sum_{i=1}^n \mathbb{P}(X_i \notin [a, b]),$$

where the last inequality is due to the union bound. By definition of a and b :

$$\begin{aligned} \mathbb{P}(A_{a,b}) &= \sum_{i=1}^n \mathbb{P}\left(\{X_i < \mu - \sigma\sqrt{n \log n}\} \cup \{X_i > \mu + \sigma\sqrt{n \log n}\}\right) \\ &= \sum_{i=1}^n \mathbb{P}\left(|X_i - \mu| > \sigma\sqrt{n \log n}\right) \leq \sum_{i=1}^n \frac{1}{n \log n} = \frac{1}{\log n}, \end{aligned}$$

where the last upper bound comes from Chebyshev's inequality. Since $\mathbb{P}(A_{a,b}) \leq \frac{1}{\log n}$ equation (5) applies, and replacing a and b gives

$$O(\mathbb{E}[\log \varepsilon]) = O\left(\log\left((b-a)\rho_f \frac{n}{K}\right)\right) = O\left(\log\left(2\sigma\sqrt{n \log n} \rho_f \frac{n}{K}\right)\right).$$

Hence, building an ESPC index with $K = n\sqrt{n \log n}$ intervals makes space usage $O(n\sqrt{n \log n})$ and guarantees that expected query time is $O(\log(2\sigma\rho_f))$. This concludes the proof. \blacktriangleleft

5 Analysis and Benchmarking

We now take a closer look at the expected time complexity. Consider an ESPC index with K subintervals for a sorted array A with n keys. Then, by Theorem 6 and Jensen's inequality:

$$\mathbb{E}[\log \varepsilon] \leq \log \mathbb{E}[\varepsilon] \leq \log\left(\frac{3n}{2} \sum_{k=1}^K p_k^2\right) = \log\left(\frac{3n}{2}\right) + \log\left(\sum_{k=1}^K p_k^2\right) \quad (6)$$

where $p_k = \mathbb{P}(X_i \in J_k)$. Recall that (p_1, \dots, p_K) constitutes an induced discrete probability distribution, describing the probabilities of the $\{X_i\}$ falling on the different $\{J_k\}$ intervals.

5.1 Analysis of Alternative Index Design

The term $\log\left(\sum_{k=1}^K p_k^2\right)$ is related to a notion of entropy and is minimized when (p_1, \dots, p_K) is a uniform distribution (see Section 5.2). Hence, we could conceivably get better performance by trying to induce $p_k = 1/K$ for all k . Consider then an alternative design where the $\{I_k\}$ subintervals are defined in such a way that they cover the $[a, b]$ range (instead of the $[X_{(1)}, X_{(n)}]$ range) and we aim to enforce $p_k = 1/K$ for all k , rather than an equal-length condition. This requires knowledge of a , b , and F , which we usually lack, but several strategies can be explored to estimate them from the data.

This new index design would refine the intervals where we expect data to be more concentrated, which seems to be a reasonable strategy. However, there is a trade-off. Choosing the subintervals so that $p_k = 1/K$ means that they will not be equal-length, unless the $\{X_i\}$ distribute uniformly. Consequently, for a given query parameter q , we can no longer find its subinterval in constant time, as we did before with $k \leftarrow \lceil \frac{1}{\delta}(q - X_{(1)}) \rceil$. This version of the index would need a new strategy to find the relevant subinterval I_k such that $q \in I_k$.

A standard strategy would be to create a new array A' formed by the left endpoints of the $\{I_k\}$, that is, $A' = [t_0, \dots, t_{k-1}]$. Now, for a given search parameter q we can find the relevant I_k with $k = \mathbf{rank}_{A'}(q)$, and then proceed with steps (5) and (6) from Algorithm 2. Notice that the index is now hierarchical, with A' at the top and A (segmented into the $\{I_k\}$) at the bottom. From equation (6), expected time at the bottom layer would scale as

$$\mathbb{E}[\log \varepsilon_{\text{bot}}] \leq \log \left(\frac{3n}{2} \sum_{k=1}^K p_k^2 \right) = \log \left(\frac{3n}{2} \sum_{k=1}^K \frac{1}{K^2} \right) = \log \left(\frac{3n}{2K} \right).$$

On the other hand, the top layer has K keys $\{t_0, \dots, t_{k-1}\}$. Since the index is designed so that $p_k = 1/K$, this means $F(t_i) = i/k$. As a consequence, the keys in A' approximately follow distribution F , with a better approximation as $K \rightarrow \infty$. Suppose we use an ESPC index with K' subintervals to speed up computation of $\mathbf{rank}_{A'}$. By Proposition 7 the expected time at the top layer will be dominated by $\mathbb{E}[\log \varepsilon_{\text{top}}] \leq \log(3(b-a)\rho_f K/(2K'))$ and the total expected time can be bounded as

$$\mathbb{E}[\log \varepsilon_{\text{bot}}] + \mathbb{E}[\log \varepsilon_{\text{top}}] \leq \log \left(\frac{3n}{2K} \right) + \log \left(\frac{3(b-a)}{2} \rho_f \frac{K}{K'} \right) = \log \left(\frac{9(b-a)}{4} \rho_f \frac{n}{K'} \right).$$

Notice that $(b-a)\rho_f$ is still present in this bound. Now, for instance, with $K = K' = n$ we get twice the space usage compared to Theorem 1 and the same asymptotic expected query time. This indicates that asymptotic complexity remains the same for a hierarchical or tree-like version of the index, so the basic ESPC index gives a simpler proof for our results.

5.2 Bounds for Expected Query Time

Complexity due to number of keys. Equation (6) provides an interesting insight. We see that the expected log-error (and hence, the expected query time) can be decomposed into two separate sources of complexity. The first is due exclusively to the number of keys n , which suggests that the query problem gets increasingly difficult when new keys are inserted, even if the underlying distribution of the data remains the same. This helps to explain why learned indexes evidence a need for updating [33, 26] in the face of new keys arriving, even when the same dataset is used to simulate the initial keys and the subsequent arrivals.

Information-theoretical considerations. The sum $\sum_{k=1}^K p_k^2$ is related to the Rényi entropy [25]. For a discrete random variable P with possible outcomes $1, \dots, K$ and correspondent probabilities $\{p_k\}_{k=1}^K$, the Rényi entropy of order α is defined as $H_\alpha(P) = \frac{1}{1-\alpha} \log \left(\sum_{k=1}^K p_k^\alpha \right)$ for $\alpha > 0, \alpha \neq 1$. The cases $\alpha = 0, 1, \infty$ are defined as limits, with $H_1(P)$ giving the Shannon entropy [25]. The Rényi entropy can be understood as a generalization of the Shannon entropy which preserves most of its properties. Consequently, equation (6) can be written as:

$$\mathbb{E}[\log \varepsilon] \leq \log \left(\frac{3n}{2} \right) - H_2(P). \quad (7)$$

H_2 appears in cryptography [4, 27] due to its relation with the collision probability $\sum_{k=1}^K p_k^2$, and is maximized by the uniform distribution [27], that is, when $p_k = 1/K$ for all k . This suggests that negative entropy might be a good measure of statistical complexity for learned indexes, beyond the ESPC index. It captures a notion of distance to the uniform distribution U , which is easy to learn for the type of simple models commonly used in learned indexes. This notion of distance can be formalized by writing $-H_2(P) = D_2(P||U) - \log K$, where $D_2(P||U)$ is the Rényi divergence of order 2, defined analogously to the KL divergence [29].

■ **Table 1** Space and expected time complexity for learned indexes and B+Tree. The probabilistic assumptions are not necessary for the index to work, but they are necessary to guarantee the stated performance. For the last three rows, it is assumed that f has bounded support $[a, b]$.

Method	Space	Expected time	Parameters / Notation	Probabilistic Model
B+Tree	$O\left(\frac{n}{b}\right)$	$O(\log n)$	Branching factor b	None
PGM [10]	$O\left(\frac{n}{\varepsilon^2}\right)$	$O(\log n)$	Maximum error ε	Gaps $\{X_{i+1} - X_i\}$ i.i.d. with finite mean μ and variance σ^2 , $\varepsilon \gg \sigma/\mu$
PCA [34]	$O(\rho^{1+\delta} n^{1+\frac{\delta}{2}})$	$O\left(\log \frac{1}{\delta}\right)$	$\delta > 0$, $\rho = (b-a)\rho_1$	$\{X_i\}$ i.i.d., density f , $f(x) \leq \rho_1 < \infty$ for all x
RDA [34]	$O(\rho n)$	$O(\log(\log n))$	$\rho = \frac{\rho_1}{\rho_2}$	Same as PCA, plus $0 < \rho_2 \leq f(x)$ for all x
ESPC	$O(n)$ $O\left(\frac{n}{\log n}\right)$	$O(\rho)$ $O(\rho + \log(\log n))$	K intervals, $\rho = \log((b-a)\rho_f)$	$\{X_i\}$ share density f , $\rho_f = \ f\ _2^2 < \infty$

Characterization via the stochastic process. A similar insight can be reached by reference to the density function f . By Proposition 7 and Jensen’s inequality, we know

$$\mathbb{E}[\log \varepsilon] \leq \log\left((b-a)\rho_f \frac{n}{K}\right) = \log n + \log(b-a) - \log(K) - h_2(X_i),$$

where we use $h_\alpha(X) = \frac{1}{1-\alpha} \log(\|f\|_\alpha^\alpha)$ to denote the Rényi entropy for continuous random variables [31], similar to how differential entropy is defined in analogy to the Shannon entropy. This last expression provides a useful bound for the expected log-error, by separating the effect of the number n of keys, the number K of subintervals (which represent space usage), the support $b-a$, and an inherent characteristic $h_2(X_i)$ of the stochastic process.

5.3 Comparison with Existing Methods

In this section, we compare our probabilistic model and complexity bounds to previous work. All results are stated in terms of our RAM model. Table 1 summarizes the main points.

Non-learned methods. There are many classical algorithms for searching a sorted array. They tend to be comparison-based and have $O(\log n)$ expected time, unless strong assumptions are made (e.g., interpolation search uses $O(\log(\log n))$ time on average if the keys are uniformly distributed [17]). We use B+Tree as a representative. A B+Tree has space overhead $O(n/b)$ where b is the branching factor. Since b does not depend on n , asymptotically this is just $O(n)$. The number of comparisons required to find a key is $O(\log n)$. This represents both worst-case and average-case complexity, with no probabilistic assumptions.

PGM-index. In the absence of probabilistic assumptions, space overhead of the PGM-index is $O(n/\varepsilon)$ and expected time is $O(\log n)$. Here, ε is a hyperparameter of the method and is independent of n , so this constitutes the same asymptotic complexity as B+Trees. Under certain assumptions, space overhead can be strengthened to $O(n/\varepsilon^2)$ [9], representing a constant factor improvement (i.e., same asymptotic space complexity).

In contrast to [34] and our work, the assumptions in [9] concern the *gaps* between keys. Specifically, the $\{X_{i+1} - X_i\}$ are assumed to be independent and identically distributed (i.i.d.), with finite mean μ and variance σ^2 . It is also assumed that $\varepsilon \gg \sigma/\mu$. This setting is

not directly comparable to ours, in the sense that neither is a particular instantiation of the other. However, the assumptions described above are very strong in our view and seem hard to justify for most practical applications. Also, the $\varepsilon \gg \sigma/\mu$ condition constrains the choice of this key hyperparameter.

PCA Index and RDA Index. In [34] several indexes are presented, the most relevant for us being the PCA and RDA indexes. The PCA Index is similar to the ESPC index, with some differences as explained in Section 3.2, including the use of binary search instead of exponential search. The RDA Index is similar to the Recursive Model Index (RMI) from [18]. All results below are stated in terms of the RAM model under the uniform cost criterion.

In [34], the same assumption regarding bounded support $[a, b]$ for f is needed, so we consider that case in this section. Extensions to unbounded support incur a penalty in terms of space complexity, dependent on the tails of the distribution. We first describe the PCA Index. For any $\delta > 0$, there is a PCA Index with space overhead $O((b-a)^{1+\delta} \rho_1^{1+\delta} n^{1+\delta/2})$ and expected time complexity $O(\log \frac{1}{\delta})$. Here, ρ_1 is an upper bound for f . Since it is independent of n , space overhead is asymptotically $O(n^{1+\delta/2})$. A direct comparison to the ESPC index is possible through Theorem 1, which provides a strictly stronger result, with expected constant time and truly linear space.

On the other hand, the RDA Index has space overhead $O(\frac{\rho_1}{\rho_2} n)$, where ρ_2 is a lower bound for f and ρ_1 is an upper bound (as for the case of the PCA Index). Asymptotically in n , this is $O(n)$. Expected query time in this case is $O(\log(\log n))$. A direct comparison to the ESPC index is possible through Theorem 2, which provides a strictly stronger result, with $O(\log(\log n))$ expected time and sublinear space overhead.

Even as we improve the complexity bounds, we also generalize the domain of application for our result. As part of the probabilistic model for the PCA and RDA indexes, it is assumed that the $\{X_i\}$ are i.i.d. with density f [34]. PCA index assumes f is upper bounded, that is, there exists some fixed $\rho_1 < \infty$ such that $f(x) \leq \rho_1$ for all x . Notice that this condition implies that f is square-integrable, so that our assumptions are strictly weaker:

$$\rho_f = \int_{\mathbb{R}} f^2(x) dx \leq \rho_1 \int_{\mathbb{R}} f(x) dx = \rho_1 < \infty.$$

The RDA index further assumes that f is bounded away from 0, that is, there exists some $\rho_2 > 0$ such that $f(x) \geq \rho_2$. Moreover, as described in [34], access to the values of a , b and ρ_1 is required to build the PCA Index, while the RDA Index requires knowledge of ρ_1 and ρ_2 . This seems like a strong assumption in practice. Finally, the analysis in [34] holds less generally than ours because it requires the $\{X_i\}$ to be independent, and it (implicitly) depends on the query parameter q having the same distribution as the $\{X_i\}$.

6 Experimental Findings

We implement the ESPC index and evaluate the results in light of the theoretical analysis.

6.1 Data

In our experiments, we use four datasets from the *Searching on Sorted Data* (SOSD) benchmark [15, 24], which has become standard in testing learned indexes [28]. Each dataset consists of a sorted array of keys, which can be used to build an index and simulate queries. Two of the datasets we use (*usparse*, *normal*) are synthetically generated. The other two

(*amzn*, *osm*) come from real-world data. In particular, *osm* is known to be challenging for learned indexes [24].

6.2 Experimental Design

Our complexity bounds (Theorems 1 and 2) are corollaries of Propositions 3 and 7, which are the key results to validate empirically. We do this in the following way. Each dataset consists of a sorted array of 20×10^7 keys. We use a subsample of $n = 10^7$ keys, so that experiments are less computationally expensive. For a fixed subsample, we take different values of K and for each:

1. We build the ESPC index with K subintervals. Space overhead is estimated by reporting the amount of memory used by the index. From Proposition 3, this should be $O(K)$. In terms of experiments, we expect a linear relationship between memory overhead and K .
2. We sample Q keys from A . We compute **rank** for these keys using the index and measure the prediction error. We then average these values to estimate the expected error. From Proposition 7, this should be smaller than $\frac{3(b-a)}{2} \rho_f \frac{n}{K}$. In practice, we use an estimate $\hat{\rho}_f$ of ρ_f (see Section 6.4) and all data is re-scaled to the $[0, 1]$ interval, so that we consider $b - a = 1$. For experiments, we plot the average prediction error alongside $(3\hat{\rho}_f n)/(2K)$, as functions of K . We expect the average error curve to be below $(3\hat{\rho}_f n)/(2K)$.

We use $Q = 30 \times 10^6$. For K we use values of $10^3, 5 \times 10^3, 10^4, 5 \times 10^4, 10^5$ and 2×10^5 .


6.3 Results

In terms of storage, our experiments show a perfect linear relationship between memory overhead and K , where **memory overhead** = $32K$ bytes. This result is in accordance to the $O(K)$ estimate in Proposition 4.1. The agreement between experiments and theory is to be expected here, because the space overhead of the ESPC index does not depend on probabilistic factors.

On the other hand, Figure 4 shows how the prediction error changes with the number of subintervals K . We plot the average experimental error along with the theoretical bound for the expected error. In accordance with Proposition 7, this theoretical bound is given by $(3n\hat{\rho}_f)/(2K)$. As expected, for all datasets this expression serves as an upper bound for the average experimental error, and both curves have a similar shape. The results show that Proposition 7 has good predictive power, even for real-world data.

6.4 Analysis of ρ_f

As seen in Figure 4, the experimental error follows the overall shape described by the theoretical bound in Proposition 7. According to the theory, the constants involved (e.g., the intercept of the curves) should depend on ρ_f , so it is crucial to have estimates for this metric. We develop a procedure that can provide an estimate $\hat{\rho}_f$ of ρ_f , based on statistical techniques to estimate the density f . Full details are available in the extended version of the paper [6]. The values computed with this method are (*usparse*, 1.20), (*normal*, 3.89), (*amzn*, 1.72), (*osm*, 32.57). The value for *osm* is an order of magnitude greater than the rest, meaning it is farther away from the uniform distribution as measured by the Rényi divergence, helping to explain why it is challenging. The dataset with the lowest value is *usparse*. Since its underlying distribution is uniform, this is in line with the analysis in Section 5.2. Figure 4 shows that higher values of $\hat{\rho}_f$ correspond to higher experimental errors.



■ **Figure 4** Average experimental error and theoretical bound for expected error, as functions of the number K of subintervals. Top row shows synthetic datasets, bottom row shows real-world datasets. The plots use a logarithmic scale for both axes, which prevents clustering of the points.

7 Conclusions

Theoretical understanding of learned indexes has not kept up with their practical development. We narrow this gap by proving that learned indexes exhibit strong theoretical guarantees under a very general probabilistic model. In particular, we prove learned indexes can achieve $O(1)$ expected query time with at most linear space. We introduce a specific index that achieves this performance, and we find a general bound for the expected time in terms of the number of keys and the Rényi entropy of the distribution. From this result, we derive a metric ρ_f that can be used to characterize the difficulty of learning the data. We describe a procedure for estimating ρ_f for real data, and we show that the complexity bounds hold in practice. Our results help explain the good experimental performance of learned indexes, under the most general conditions considered so far in the literature. As future work, it is theoretically important to prove lower bounds for expected query time. We believe this type of result may hinge on information-theoretical properties (e.g., through rate-distortion theory [3]), which our analysis shows are relevant for learned indexes. On the other hand, it is also important to extend our analysis to the case of dynamic datasets, where the probability distribution may exhibit drift over time and updates are necessary. This is an important case in applications, and has not been considered in theoretical analysis.

References

- 1 Alfred V. Aho and John E. Hopcroft. *The Design and Analysis of Computer Algorithms*. Addison-Wesley Longman Publishing Co., Inc., USA, 1st edition, 1974.
- 2 Jon Louis Bentley and Andrew Chi-Chih Yao. An almost optimal algorithm for unbounded searching. *Information Processing Letters*, 5(3):82–87, 1976. doi:10.1016/0020-0190(76)90071-5.
- 3 Toby Berger. Rate-distortion theory. *Wiley Encyclopedia of Telecommunications*, 2003.
- 4 Christian Cachin. *Entropy measures and unconditional security in cryptography*. PhD thesis, ETH Zurich, 1997. URL: <https://d-nb.info/950686247>.
- 5 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, Cambridge, MA, USA, 3rd edition, 2009.

- 6 Luis Croquevielle, Guang Yang, Liang Lian, Ali Hadian, and Thomas Heinis. Querying in constant expected time with learned indexes. *arXiv preprint arXiv:2405.03851*, 2024.
- 7 H.A. David and H.N. Nagaraja. *Order Statistics*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2004.
- 8 Jialin Ding, Umar Farooq Minhas, Jia Yu, Chi Wang, Jaeyoung Do, Yinan Li, Hantian Zhang, Badrish Chandramouli, Johannes Gehrke, Donald Kossmann, David Lomet, and Tim Kraska. Alex: An updatable adaptive learned index. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, SIGMOD '20, pages 969–984, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3318464.3389711.
- 9 Paolo Ferragina, Fabrizio Lillo, and Giorgio Vinciguerra. Why are learned indexes so effective? In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3123–3132, Virtual, 13–18 July 2020. PMLR. URL: <http://proceedings.mlr.press/v119/ferragina20a.html>.
- 10 Paolo Ferragina and Giorgio Vinciguerra. The PGM-index: A fully-dynamic compressed learned index with provable worst-case bounds. *Proc. VLDB Endow.*, 13(8):1162–1175, April 2020. doi:10.14778/3389133.3389135.
- 11 Alex Galakatos, Michael Markovitch, Carsten Binnig, Rodrigo Fonseca, and Tim Kraska. Fiting-tree: A data-aware index structure. In *Proceedings of the 2019 International Conference on Management of Data*, SIGMOD '19, pages 1189–1206, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3299869.3319860.
- 12 Rodrigo González, Szymon Grabowski, Veli Mäkinen, and Gonzalo Navarro. Practical implementation of rank and select queries. In *Poster Proc. Volume of 4th Workshop on Efficient and Experimental Algorithms (WEA)*, pages 27–38, Athens, Greece, 2005. CTI Press and Ellinika Grammata.
- 13 Goetz Graefe et al. Modern B-tree techniques. *Foundations and Trends in Databases*, 3(4):203–402, 2011. doi:10.1561/19000000028.
- 14 Ali Hadian and Thomas Heinis. Shift-table: A low-latency learned index for range queries using model correction. In Yannis Velegrakis, Demetris Zeinalipour-Yazti, Panos K. Chrysanthis, and Francesco Guerra, editors, *Proceedings of the 24th International Conference on Extending Database Technology, EDBT 2021, Nicosia, Cyprus, March 23 - 26, 2021*, pages 253–264. OpenProceedings.org, 2021. doi:10.5441/002/edbt.2021.23.
- 15 Andreas Kipf, Ryan Marcus, Alexander van Renen, Mihail Stoian, Alfons Kemper, Tim Kraska, and Thomas Neumann. SOSD: A benchmark for learned indexes. *arXiv preprint arXiv:1911.13014*, 2019. arXiv:1911.13014.
- 16 Andreas Kipf, Ryan Marcus, Alexander van Renen, Mihail Stoian, Alfons Kemper, Tim Kraska, and Thomas Neumann. Radixspline: A single-pass learned index. In *Proceedings of the Third International Workshop on Exploiting Artificial Intelligence Techniques for Data Management*, aiDM '20, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3401071.3401659.
- 17 Donald E. Knuth. *The Art of Computer Programming, Volume 3: (2nd Ed.) Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., USA, 1998.
- 18 Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, SIGMOD '18, pages 489–504, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3183713.3196909.
- 19 Justin J. Levandoski, David B. Lomet, and Sudipta Sengupta. The bw-tree: A b-tree for new hardware platforms. In *Proceedings of the 2013 IEEE International Conference on Data Engineering (ICDE 2013)*, ICDE '13, pages 302–313, USA, 2013. IEEE Computer Society. doi:10.1109/ICDE.2013.6544834.
- 20 Ling Liu and M Tamer Özsu. *Encyclopedia of Database Systems*, volume 6. Springer New York, NY, USA, 2009.

- 21 Scott Lloyd and Maya Gokhale. Near memory key/value lookup acceleration. In *Proceedings of the International Symposium on Memory Systems, MEMSYS '17*, pages 26–33, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3132402.3132434.
- 22 Zhanyu Ma and Arne Leijon. Coding bounded support data with beta distribution. In *2010 2nd IEEE International Conference on Network Infrastructure and Digital Content*, pages 246–250, 2010. doi:10.1109/ICNIDC.2010.5657779.
- 23 Zhanyu Ma, Andrew E. Teschendorff, Arne Leijon, Yuanyuan Qiao, Honggang Zhang, and Jun Guo. Variational bayesian matrix factorization for bounded support data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(4):876–889, 2015. doi:10.1109/TPAMI.2014.2353639.
- 24 Ryan Marcus, Andreas Kipf, Alexander van Renen, Mihail Stoian, Sanchit Misra, Alfons Kemper, Thomas Neumann, and Tim Kraska. Benchmarking learned indexes. *Proc. VLDB Endow.*, 14(1):1–13, September 2020. doi:10.14778/3421424.3421425.
- 25 Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, volume 4, pages 547–562, Oakland, CA, USA, 1961. University of California Press.
- 26 Amirhesam Shahvarani and Hans-Arno Jacobsen. Parallel index-based stream join on a multicore cpu. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data, SIGMOD '20*, pages 2523–2537, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3318464.3380576.
- 27 Maciej Skórski. Shannon entropy versus renyi entropy from a cryptographic viewpoint. In Jens Groth, editor, *Cryptography and Coding*, pages 257–274, Cham, 2015. Springer International Publishing. doi:10.1007/978-3-319-27239-9_16.
- 28 Zhaoyan Sun, Xuanhe Zhou, and Guoliang Li. Learned index: A comprehensive experimental evaluation. *Proc. VLDB Endow.*, 16(8):1992–2004, April 2023. doi:10.14778/3594512.3594528.
- 29 Tim van Erven and Peter Harremo. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014. doi:10.1109/TIT.2014.2320500.
- 30 Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, UK, 2018.
- 31 Susana Vinga and Jonas S Almeida. Rényi continuous entropy of dna sequences. *Journal of theoretical biology*, 231(3):377–388, 2004.
- 32 Chaichon Wongkham, Baotong Lu, Chris Liu, Zhicong Zhong, Eric Lo, and Tianzheng Wang. Are updatable learned indexes ready? *Proc. VLDB Endow.*, 15(11):3004–3017, July 2022. doi:10.14778/3551793.3551848.
- 33 Guang Yang, Liang Liang, Ali Hadian, and Thomas Heinis. FLIRT: A fast learned index for rolling time frames. In Julia Stoyanovich, Jens Teubner, Nikos Mamoulis, Evaggelia Pitoura, Jan Mühlig, Katja Hose, Sourav S. Bhowmick, and Matteo Lissandrini, editors, *Proceedings 26th International Conference on Extending Database Technology, EDBT 2023, Ioannina, Greece, March 28-31, 2023*, pages 234–246, Ioannina, Greece, 2023. OpenProceedings.org. doi:10.48786/edbt.2023.19.
- 34 Sepanta Zeighami and Cyrus Shahabi. On distribution dependent sub-logarithmic query time of learned indexing. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org, 2023.

A Proof of Proposition 8

Proof of Proposition 8. The proof follows closely the argument used to prove Theorem 6 and Proposition 7. By the tower property for conditional expectation, $\mathbb{E}[\varepsilon] = \mathbb{E}[\mathbb{E}[\varepsilon \mid X_1, \dots, X_n]]$,

19:20 Querying in Constant Expected Time with Learned Indexes

where the inner expected value gives

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] = \int_{\mathbb{R}} \varepsilon(q)g(q)dq = \int_{X_{(1)}}^{X_{(n)}} \varepsilon(q)g(q)dq.$$

where the last equality comes from the fact that for $q \notin [X_{(1)}, X_{(n)}]$ we know the exact value of $\mathbf{rank}(q)$, which is either 0 or n , and hence we have that $\varepsilon(q) = 0$. Replacing $\varepsilon(q) = |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|$ and noticing that the $\{I_k\}$ form a partition of $[X_{(1)}, X_{(n)}]$:

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] = \sum_{k=1}^K \int_{I_k} |\mathbf{rank}(q) - \hat{\mathbf{r}}(q)|g(q)dq.$$

By Lemma 4 we can bound the error within each I_k :

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] \leq \frac{1}{2} \sum_{k=1}^K n_k \int_{I_k} g(q)dq,$$

where n_k is the number of keys in subinterval I_k . Now, denote by $n_{k,\ell}$ the number of keys in $I_k \cap J_\ell$. Then, we can write

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] \leq \frac{1}{2} \sum_{k=1}^K \left(\sum_{\ell=1}^K n_{k,\ell} \right) \int_{I_k} g(q)dq = \frac{1}{2} \sum_{\ell=1}^K \sum_{k=1}^K n_{k,\ell} \int_{I_k} g(q)dq, \quad (8)$$

where first we have written n_k as $\sum_{\ell=1}^K n_{k,\ell}$ and then we have exchanged the order of summation. Now, denote by $L(k)$ the set of ℓ indices such that $I_k \cap J_\ell \neq \emptyset$. Notice that $n_{k,\ell} = 0$ if $\ell \notin L(k)$. Hence, inequality (8) becomes:

$$\mathbb{E}[\varepsilon \mid X_1, \dots, X_n] \leq \frac{1}{2} \sum_{\ell=1}^K \sum_{k:\ell \in L(k)} n_{k,\ell} \int_{I_k} g(q)dq. \quad (9)$$

By Fact 5 we know that $I_k \subseteq J_{\ell-1} \cup J_\ell \cup J_{\ell+1}$ for all ℓ and k such that $\ell \in L(k)$. Now, define q_ℓ as $\mathbb{P}(q \in J_\ell)$ for each $\ell = 1, \dots, K$, analogously to how the p_ℓ are defined as $\mathbb{P}(X_i \in J_\ell)$. Using this notation, from inequality (9) we get

$$\begin{aligned} \mathbb{E}[\varepsilon \mid X_1, \dots, X_n] &\leq \frac{1}{2} \sum_{\ell=1}^K \sum_{k:\ell \in L(k)} n_{k,\ell} \left(\int_{J_{\ell-1}} g(q)dq + \int_{J_\ell} g(q)dq + \int_{J_{\ell+1}} g(q)dq \right) \\ &= \frac{1}{2} \sum_{\ell=1}^K (q_{\ell-1} + q_\ell + q_{\ell+1}) \sum_{k:\ell \in L(k)} n_{k,\ell} \\ &= \frac{1}{2} \sum_{\ell=1}^K (q_{\ell-1} + q_\ell + q_{\ell+1}) m_\ell, \end{aligned} \quad (10)$$

where m_ℓ denotes the number of keys in J_ℓ . Now, applying expected value on both sides of inequality (10), we get that

$$\mathbb{E}[\varepsilon] = \mathbb{E}[\mathbb{E}[\varepsilon \mid X_1, \dots, X_n]] \leq \frac{1}{2} \sum_{\ell=1}^K (q_{\ell-1} + q_\ell + q_{\ell+1}) \mathbb{E}[m_\ell] \quad (11)$$

Furthermore, we have

$$\mathbb{E}[m_\ell] = \mathbb{E} \left[\sum_{i=1}^n \mathbf{1}_{X_i \in J_\ell} \right] = \sum_{i=1}^n \mathbb{E}[\mathbf{1}_{X_i \in J_\ell}] = \sum_{i=1}^n \mathbb{P}(X_i \in J_\ell) = np_\ell.$$

Replacing $\mathbb{E}[m_\ell] = np_\ell$ in inequality (11), we get:

$$\mathbb{E}[\varepsilon] \leq \frac{1}{2} \sum_{\ell=1}^K (q_{\ell-1} + q_\ell + q_{\ell+1}) \mathbb{E}[m_\ell] = \frac{n}{2} \sum_{\ell=1}^K (q_{\ell-1}p_\ell + q_\ell p_\ell + q_{\ell+1}p_\ell). \quad (12)$$

Now, by Cauchy-Schwarz inequality we know that

$$\sum_{\ell=1}^K q_{\ell-1}p_\ell \leq \left(\sum_{\ell=1}^K q_{\ell-1}^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2} = \left(\sum_{\ell=0}^{K-1} q_\ell^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2} \leq \left(\sum_{\ell=1}^K q_\ell^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2},$$

where we have used the fact that $q_0 = \mathbb{P}(q \in J_0) = 0$ since J_0 is defined as the empty set. Similarly, by Cauchy-Schwarz inequality and the fact that $q_{K+1} = 0$ we can prove

$$\sum_{\ell=1}^K q_\ell p_\ell \leq \left(\sum_{\ell=1}^K q_\ell^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2}, \quad \sum_{\ell=1}^K q_{\ell+1}p_\ell \leq \left(\sum_{\ell=1}^K q_\ell^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2}.$$

Replacing in (12) we get the new inequality

$$\mathbb{E}[\varepsilon] \leq \frac{3n}{2} \left(\sum_{\ell=1}^K q_\ell^2 \right)^{1/2} \left(\sum_{\ell=1}^K p_\ell^2 \right)^{1/2}.$$

Now, since both f and g are square-integrable functions, the same argument from Proposition 7 can be used to prove that

$$\sum_{\ell=1}^K q_\ell^2 \leq \frac{(b-a)}{K} \rho_g, \quad \text{and} \quad \sum_{\ell=1}^K p_\ell^2 \leq \frac{(b-a)}{K} \rho_f.$$

Putting everything together:

$$\mathbb{E}[\varepsilon] \leq \frac{3n}{2} \left(\sum_{\ell=1}^K q_\ell^2 \sum_{\ell=1}^K p_\ell^2 \right)^{1/2} \leq \frac{3n}{2} \left(\frac{(b-a)}{K} \rho_g \frac{(b-a)}{K} \rho_f \right)^{1/2} = \frac{3(b-a)}{2} \sqrt{\rho_f \rho_g} \frac{n}{K}.$$

This concludes the proof. ◀