
LISA: Towards Learned DNA Sequence Search

Darryl Ho*, Jialin Ding*, Sanchit Misra†, Nesime Tatbul*†,
Vikram Nathan*, Vasimuddin Md†, Tim Kraska*
*Massachusetts Institute of Technology, †Intel Labs

Abstract

Next generation sequencing (NGS) technologies have enabled affordable sequencing of billions of short DNA fragments at high throughput, paving the way for population-scale genomics. Genomics data analytics at this scale requires overcoming performance bottlenecks, such as searching for short DNA sequences over long reference sequences. In this paper, we introduce LISA (Learned Indexes for Sequence Analysis), a novel learning-based approach to DNA sequence search. As a first proof of concept, we focus on accelerating one of the most essential flavors of the problem, called *exact search*. LISA builds on and extends FM-index, which is the state-of-the-art technique widely deployed in genomics tool-chains. Initial experiments with human genome datasets indicate that LISA achieves up to a factor of $4\times$ performance speedup against its traditional counterpart.

1 Introduction

Rapid advances in high-throughput next generation sequencing (NGS) technologies have enabled affordable sequencing of billions of short DNA fragments (called “reads”) at an unprecedented scale. For example, the Illumina NovaSeq 6000 Sequencer can sequence 20 billion reads of length 150 each in less than 2 days, generating 6 Terabases of data at a low cost of about \$1000 per human genome [1]. Already today, a growing number of public and private sequencing centers with hundreds of NGS deployments pave the way for population-level genomics. However, realizing this vision in practice heavily relies on building scalable systems for high-speed genomics data analysis.

DNA sequence alignment plays a critical role in genome analysis. In its simplest form, an aligner tries to piece together the short reads by mapping each individual read to a long reference genome (e.g., the human genome consisting of 3 billion bases). The key operation that has been shown to constitute a significant performance bottleneck during this mapping process is the search for *exact* or *inexact* matches of read substrings over the given reference sequence [2–7]. In this paper, we focus on the exact-search variant of this search problem.

The state-of-the-art technique to perform exact search is based on building an FM-index over the reference genome [8]. The key idea behind an FM-index is that in the lexicographically sorted order of all suffixes of the reference sequence, all matches of a short DNA sequence (a.k.a., a “query”) will fall in a single region matching the prefixes of contiguously located suffixes. Over the years, many improvements have been made to make the FM-index more efficient, leading to several state-of-the-art implementations that are highly cache- and processor-optimized [3–7, 9–15, 10, 15]. Hence, it becomes increasingly more challenging to further improve this critical step in the genomics pipeline to scale with increasing data growth.

In this paper, we propose a different approach to improving the sequence search performance: LISA (Learned Indexes for Sequence Analysis). The core idea behind LISA, which enables a new ML-enhanced algorithm for DNA search, is to speed up the process of finding the right region of suffixes by learning the distribution of suffixes in the reference. We do this in a way similar to how learned indexes capture data value distributions through models learned from data [16].

When evaluated on an Intel® Core™ i9-9900K 3.6 GHz processor, despite being single-threaded and not yet fully optimized to the underlying hardware architecture, our current implementation achieves

nearly $4\times$ speedup against a state-of-the-art single-threaded, CPU-optimized version of the FM-index based algorithm [15], for a workload of 50 million queries matched against the human genome. This early result shows that learned DNA sequence search is a promising idea.

To the best of our knowledge, this is the first work exploring how ML-enhanced algorithms can improve the process of DNA sequence search, while providing identical semantic guarantees as the traditional algorithms. In summary, this paper makes the following contributions:

- enhancements to the FM-index that enable the application of learning-based search,
- a new search algorithm that applies the learning-based approach to the enhanced FM-index to find all exact matches,
- an experimental comparison of our approach against a highly-tuned baseline using realistic workloads on the human genome.

In the rest of this paper, we first provide background on the traditional FM-index based exact search algorithm and the idea of learned index structures; then we present our new approach LISA, along with results from our experimental study. A more detailed technical report can be found at [17].

2 Background

A DNA sequence is a string over the alphabet, $\Sigma = \{A, C, G, T\}$, representing the four bases. For the rest of this paper, we use the terms “sequence” and “string” interchangeably, as well as the terms “base” and “character”. Exact DNA sequence search is a key kernel in many genomics tools, including the widely-used sequence mapping tool Bowtie 2 [4]. Given a reference sequence R and a query sequence Q , the goal of exact sequence search is to find exact end-to-end matches of Q in R . Typically, $|R| \approx 10^9$ bases; e.g., the length of the human genome is around 3 billion bases. On the other hand, $|Q|$ is typically less than 200 bases; e.g., the default query length in Bowtie 2 is 21.

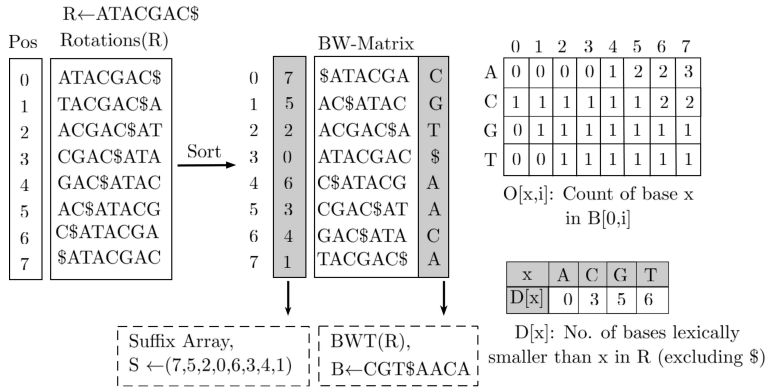


Figure 1: FM-index (S , B , D , O) and BW-matrix for sample reference sequence $R \leftarrow \text{ATACGAC}\$$. The lexicographical ordering is $\$ < A < C < G < T$ [18].

FM-index: Fig. 1 depicts the construction of the FM-index for an example reference sequence R . First, we append R with the character $\$ \notin \Sigma$ which is lexicographically smaller than all characters in Σ . Subsequently, we obtain all the rotations of R (Rotations(R)). The lexicographically sorted order of the rotations forms the BW-matrix. The BWT (B) is the last column of the BW-matrix. The original positions in R of the first bases of these rotations constitute the suffix array (S).

All the exact matches of a query can be found as prefixes of the rotations in the BW-matrix. Since the BW-matrix is lexicographically sorted, these matches are located in contiguous rows of the BW-matrix. Therefore, for a query, all the matches can be represented as a range of rows of the BW-matrix. This range is called the *SA interval* of the query. For example, in Fig. 1, the *SA interval* of query “AC” is [1,2]. The values of the suffix array in the SA interval are 5 and 2. Indeed, the sequence “AC” is found at positions 5 and 2 in the reference sequence.

The FM-index is used to expedite search for the SA interval [19]. It consists of the suffix array S and the BWT B , as well as D and O data structures. $D(x)$ is the count of bases in $R[0, |R| - 1]$ that are lexicographically smaller than $x \in \Sigma$. $O(x, i)$ is the count of occurrences of base x in $B[0, i]$. Note that the BW-matrix is not stored.

Using the FM-index, finding the SA interval is done using the *backward-search algorithm*. For a reference sequence of length n , the algorithm initializes the SA interval to $[0, n)$. The SA interval is updated over the course of the algorithm. Specifically, the algorithm processes the query sequence backwards, starting from the last character, prepending one character at a time and updating the SA interval after each prepended character in $O(1)$ time. To update the SA interval, the FM-index supports a function $f : (char, int) \rightarrow int$ that takes the prepended character c and an integer location i , and finds in $O(1)$ time the lower-bound location i' in the BW-matrix of string $\text{concat}(c, \text{BW-matrix}[i])$. (The lower bound of a string s is the first entry in the BW-matrix which does not compare less than s .) Given an SA interval $[l, u)$, after prepending the character c , the SA interval is updated to $[f(c, l), f(c, u))$. Since prepending each character takes $O(1)$ time, the overall algorithm takes $O(|Q|)$ time. For further information on the FM-index, see [8].

Learned Indexes: Recent work on learned index structures has introduced the idea that indexes are essentially models that map inputs to positions and, therefore, can be replaced by other types of models, such as machine learning models [16]. For example, a B-tree index maps a given key to the position of that key in a sorted array. Kraska et al. show that using knowledge of the distribution of keys can produce a learned model, called the recursive model index (RMI), that outperforms B-trees in query time and memory footprint. Following [16], other learned index structures have been proposed [20–23].

Taking a similar perspective, the FM-index can be seen as a model that maps a given query sequence to the SA interval for that query sequence. Based on this insight, in LISA, we use knowledge of the distribution of subsequences within the reference sequence to create a learned index structure that enables faster exact search queries.

3 LISA

Backward-search algorithm performs exact search of a query Q using the FM-index by iterating through the query sequence in backwards order, one base at a time, thereby consuming $O(|Q|)$ time. The key idea of LISA is to iterate backwards through the query sequence in chunks of K bases at a time, so that exact search takes $O(|Q|/K)$ time. To do this, LISA requires two components: (1) a new data structure called the IP-BWT that enables processing K bases at a time, and (2) a method to efficiently search through the IP-BWT, for which we use an RMI. Similar to the *backward-search algorithm*, the goal of LISA is to output the SA interval of a query sequence.

IP-BWT: In the *backward-search algorithm*, in order to update the SA interval after prepending a character, FM-index supports a function $f : (char, int) \rightarrow int$ that takes the character c and an integer location i , and finds in $O(1)$ time the lower-bound location i' in the BW-matrix of string $\text{concat}(c, \text{BW-matrix}[i])$. For LISA, we need to support a similar function that takes in a length- K string s and a location i , and returns another location i' (i.e. $f_K : (char^K, int) \rightarrow int$) representing the lower-bound location in the BW-matrix of string $\text{concat}(s, \text{BW-matrix}[i])$.

For this purpose, we introduce the Index-Paired BWT (IP-BWT) array. Each entry of IP-BWT consists of a $(char^K, int)$ pair. The first part is the first K characters of the corresponding BW-matrix row. The second part is the BW-matrix location of the string with the first K and the last $n - K$ characters swapped. Our desired function $f_K : (char^K, int) \rightarrow int$ is now equivalent to finding the lower-bound location of the input $(char^K, int)$ pair in the IP-BWT array. We are free to choose any implementation for how to find that lower-bound location; for example, since the IP-BWT is sorted, we could do a binary search over the entries of the IP-BWT. Fig. 2 shows how to create an IP-BWT with $K = 3$.

BW-Matrix		IP-BWT	
00	\$CATTATTAGGA	00	\$CA, 11
01	A\$CATTATTAGG	01	A\$C, 04
02	AGGA\$CATTATT	02	AGG, 01
03	ATTAGG\$CATT	03	ATT, 02
04	ATTATTAGG\$C	04	ATT, 03
05	CATTATTAGG\$	05	CAT, 09
06	G\$CATTATTAG	06	G\$, 05
07	GG\$CATTATTA	07	GGA, 00
08	TAGG\$CATTAT	08	TAG, 06
09	TATTAGG\$CAT	09	TAT, 08
10	TTAGG\$CATT	10	TTA, 07
11	TTATTAGG\$CA	11	TTA, 10

Figure 2: Conversion of a BW-Matrix to IP-BWT on reference CATTATTAGGA, where $K = 3$.

Alg. 1 shows the exact search algorithm using IP-BWT. For example, using the reference sequence and IP-BWT from Fig. 2, let the query sequence be ATTA. We split this into two chunks: ATT and A (line 2). We first use the RMI to find the lower bound locations of $(A\$A, 0)$ and (ATT, n) , which are 1 and 5, respectively. We then use the RMI to find the lower bound locations of $(ATT, 1)$ and $(ATT, 5)$, which are 3 and 5. Our algorithm gives the interval $[3, 5)$. We can confirm that ATTA can be found in position 3 and 4 of the BW-matrix.

Algorithm 1 Exact Search Algorithm using IP-BWT

Input: Q , a query string; $f_K : (\text{char}^K, \text{int}) \rightarrow \text{int}$, a function that finds the lower-bound location of $(\text{char}^K, \text{int})$ in the IP-BWT

Output: $[low, high)$, an SA interval.

```
1:  $low, high \leftarrow 0, n$ 
2: split  $Q$  into  $\lceil |Q|/K \rceil$  chunks, each of length  $K$ , with the final chunk possibly shorter than  $K$ 
3: for  $\mathcal{C}$  in reversed order of chunks do
4:   if  $|\mathcal{C}| < K$  then
5:     /* Special case for when final chunk has length less than  $K$  */
6:      $C_{low} \leftarrow \mathcal{C} + "\$"$  + "A"  $\times (K - |\mathcal{C}| - 1)$ 
7:      $C_{high} \leftarrow \mathcal{C} + "T"$   $\times (K - |\mathcal{C}|)$ 
8:      $low, high \leftarrow f_K((C_{low}, low)), f_K((C_{high}, high))$ 
9:   else
10:     $low, high \leftarrow f_K((\mathcal{C}, low)), f_K((\mathcal{C}, high))$ 
11:  end if
12: end for
13: return  $[low, high)$ 
```

Faster Chunk Processing using RMI: Using the IP-BWT, we are able to process the query sequence in chunks of K bases at a time. However, when processing each chunk, we must evaluate the function f_K , which involves a binary search over the IP-BWT. This takes $O(\log n)$ time, where n is the number of entries in the IP-BWT, which is equivalent to the length of the reference sequence. Therefore, the overall runtime of exact search using IP-BWT and evaluating f_K with binary search is $O(|Q|\log n/K)$. For large reference sequences, this might be slower than *backward-search* using the FM-index.

In order to avoid paying the cost of a binary search for each evaluation of f_K , we use a learned approach to support $O(1)$ evaluation of f_K . In particular, f_K is a model that maps input keys $((\text{char}^K, \text{int})$ pairs) to their positions in the sorted IP-BWT. We model f_K using the RMI, which is a hierarchy of models that is quick to evaluate [16]; the RMI conceptually resembles a hierarchical mixture of experts [24]. Fig. 3 shows an example of using a 3-layer RMI to evaluate f_K in three steps: (1) since the RMI only accepts numbers as inputs, we first convert the input $(\text{char}^K, \text{int})$ into a number. Since the alphabet Σ only has 4 characters, any character can be represented in 2 bits. Therefore, we convert char^K into a number with $2K$ bits by concatenating the bits of the individual characters together. We then append the bits of the int . Note that we have a special case for handling the sentinel character $\$$ while maintaining this 2-bit encoding. (2) We give the encoded input to the RMI and traverse down the layers of the RMI to a leaf model. The leaf model predicts the position in the IP-BWT where it expects to find the input pair. (3) If the predicted position does not contain the input pair, we use linear search over the IP-BWT starting from the predicted position to find the actual position of the pair.

Optimization for Batched Queries: So far, we have presented our method of processing a single query. However, in practice, queries arrive in large batches. For large query batches, we can adopt an additional optimization. Because the final layer of the RMI (i.e., RMI leaves), which are responsible for predicting locations, are arranged in sorted order, if the queries also arrive in sorted order, we can use the *double-pointer technique* to find for each query its corresponding RMI leaf. This resembles the merge step of merge-sort. When the number of queries is $\Omega(\text{number of RMI leaves})$, the amortized time for each query to find its RMI leaf is $O(1)$. This avoids the cost of traversing down the RMI for each query, and also improves cache locality. We defer the details of this algorithmic improvement to our extended technical report to be released on arXiv.

Discussion: LISA’s speed advantage over FM-index comes from two components: (1) the IP-BWT, which allows LISA to process K -character chunks of the query at a time, whereas FM-index processes one character at a time, and (2) using an RMI to process each K -character chunk in $O(1)$ time, whereas a naive binary search would take $O(\log n)$ time per chunk. Therefore, an exact search query using FM-index takes $O(|Q|)$ time, whereas LISA takes $O(|Q|/K)$ time.

To discuss memory consumption, we measure in bytes in terms of n , the length of the reference sequence. The reference sequence itself therefore takes $0.25n$ space. In the state-of-the-art implementation of *backward-search algorithm*, the space of suffix array is $4n$, a compressed structure that combines BWT and O is $2n$, and D is negligible, for a total space consumption of $6n$. For LISA, the space of the suffix array is $4n$, IP-BWT is $(0.25K + 4)n$, and the RMI is usually around $0.5n$, for a total space

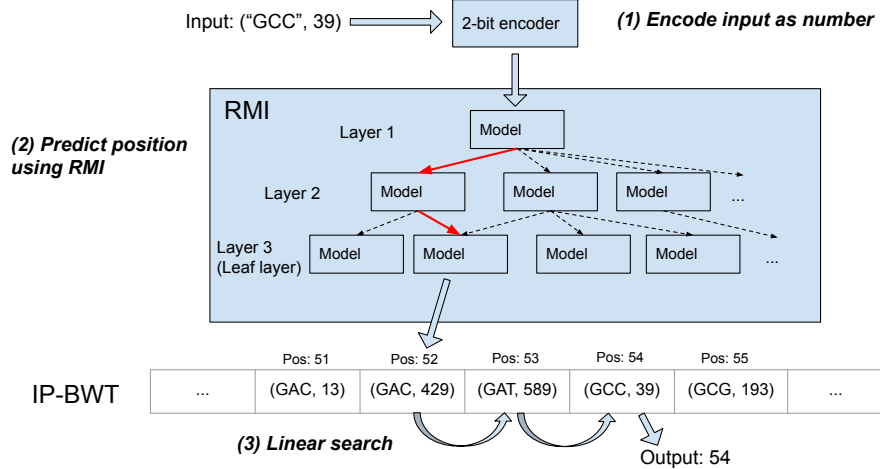


Figure 3: Using a 3-layer RMI to evaluate f for an example input $(GCC, 39)$. The RMI predicts position 52, and linear search finds the correct position. The solid red lines show the traversal path down the RMI.

consumption of $8.5n + 0.25Kn$. For example, in Bowtie 2 the default value of $|Q|$ is 21, so using an IP-BWT with $K = 21$, LISA takes around $13.75n$ space, which is larger than the FM-index. However, slightly larger space consumption is usually not a concern in practice, and if necessary LISA can use smaller K or compress the IP-BWT. Though the space usage of LISA would increase with larger values of K , we find through experiments that LISA maintains good performance using an IP-BWT with $K = 21$, even for large query lengths. Therefore, the space of LISA does not need to grow beyond around $13.75n$.

Could we replace the RMI with some other index structure that can evaluate f_K over the IP-BWT even faster? [16] showed that RMIs perform better than binary search and B trees. A lookup table (implemented as an array) that stores the output of f_K for every possible pair $(char^K, int)$ would also allow $O(1)$ evaluations but would far exceed memory capacity. However, it is possible to combine a downsampled version of the lookup table with binary search, which we discuss in the evaluation. Another idea is to use a hash table to map all n existing pairs $(char^K, int)$ to their positions; however, this would fail when we need to evaluate f_K on pairs that do not exist, such as when processing a chunk with length less than K where we pad additional characters. Also, note that since the entries of the IP-BWT must be sorted in order to maintain one contiguous SA interval, the IP-BWT itself cannot be replaced with a hash table to enable faster searches.

4 Evaluation

We present preliminary results for LISA. Experiments use SIMD, running a single-thread implementation on an Ubuntu system with Intel® Core™ i9-9900K 3.6GHz processor and 64GB RAM. As our baseline, we use a highly CPU-optimized implementation of backward search algorithm that is significantly faster than its alternatives [15]. We also compare with doing backwards search using IP-BWT and binary search (i.e., without the RMI).

Workload Scenario: We evaluate LISA on a real-world scenario: we use whole human genome as the reference sequence and for query sets, we use large sets of short query sequences of various lengths as would be found in several prominent sequence aligners. We train LISA using an IP-BWT with $K = 21$. The RMI has three layers, only uses linear regression models, and the number of models at each layer is tuned once for optimal performance on the reference sequence; the RMI is not re-trained or re-tuned for each experiment. We evaluate on four different query sequence lengths: (1) $|Q| = 21$, so that LISA processes exactly one chunk, (2) $|Q| = 42$, so that LISA processes multiple whole chunks, (3) $|Q| = 32$, so that LISA processes a chunk shorter than K , and (4) a very long query, $|Q| = 200$. For each of these query sequence lengths, we generate a batch of 50 million query sequences randomly from the human genome to be our query set, and run them together with LISA, using the batched-query optimization from §3.

Tab. 1 shows that LISA is $2.73\times$ to $3.97\times$ faster than the optimized FM-index baseline on these query lengths. Query lengths that are a perfect multiple of K (i.e., 21 and 42) perform the best. The query length of 32 has slightly lower relative performance, as LISA must still process 2 chunks per query, as if the query has length 42. The very long query sequence has lower relative performance, as LISA’s performance does not scale linearly: for longer queries, the time spent on sorting grows super-linearly. LISA achieves around a $2\times$ performance boost from using the RMI instead of binary search.

	$ Q =21$	$ Q =32$	$ Q =42$	$ Q =200$
Optimized FM-index	1509	2414	3284	15411
IP-BWT with binary search	785 (1.92 \times)	1424 (1.70 \times)	1779 (1.85 \times)	10414 (1.48 \times)
LISA	383 (3.94 \times)	762 (3.17 \times)	827 (3.97 \times)	5646 (2.73 \times)

Table 1: Average query time (measured in ticks per query) of exact search on 50 million queries, while varying query length. Relative speedup to FM-index is shown in parentheses.

Batch Size: In order to measure the effect of batch size on query times, we fix the query sequence length to $|Q|=21$ and vary the batch size from 1 to 500 million. Fig. 4 shows that as batch size increases, LISA’s time per query decreases. For batches larger than 10 thousand, LISA starts to outperform the FM-index. LISA’s performance is poor for small batch sizes, where the double pointer technique actually hurts performance, as the merge step will skip over many leaves. It is possible to achieve comparable/better performance than FM-index for small batch sizes by not using the double-pointer technique. However, in real-world scenarios, it is rare to have such small batch sizes where this would matter.




Figure 4: LISA’s performance advantage increases with batch size.

Alternative to RMI: To evaluate the effectiveness of RMI against alternative index structures, we compared LISA to a version which evaluates f_K using a combination of a lookup-table and binary search, instead of using the RMI. If we treat a $(char^K, int)$ pair as a $(2K + 32)$ -bit number, we maintain a “downsampled” lookup-table (implemented as an array) with 2^p entries, where $p < 2K + 32$. Let $q = 2K + 32 - p$. The i -th entry of the lookup-table holds the lower-bound location in the IP-BWT of the pair represented by the number $i \cdot q$. Essentially, our downsampled lookup-table holds every q -th entry of a “full” lookup-table that contains every possible $(char^K, int)$ pair. To evaluate f_K on an input represented as a $2K + 32$ -bit number, we do a lookup in the lookup-table using the first p bits to obtain lower and upper bounds on the location of the input pair in the IP-BWT, then do a binary search with the remaining q bits on the IP-BWT starting from those bounds to find the location of the input. This is our best effort at an alternative index structure that is most competitive with RMI.

For $K = 21$, $|Q| = 21$, and batch size of 50 million, replacing the RMI with a downsampled lookup-table of the same memory size results in 40% slower query times. More performant lookup-tables use significantly more memory: the lookup-table itself has space 2^{p+2} bytes, so even with only $p = 34$, the lookup table takes 69GB, which exceeds the memory of our machine. Therefore, the lookup-table is impractical and we do not include it in this paper.

5 Conclusions

In this work, we introduced our learned indexing approach for DNA sequence search, LISA, which produces promising results for the exact search problem when tested on workloads of realistic queries against the human genome. In particular, LISA achieves up to nearly $4\times$ faster query times than an extensively optimized version of the FM-index based approach, which has been the common practice in sequence search. We believe that the core ideas behind LISA can be extended to other types of DNA sequence search problems. In particular, we are currently working on using learned indexes to speed up searching for super maximal exact matches (SMEMs) for a query in the reference. For any position in query, an SMEM is the longest substring of the query through that position that has an exact match in the reference [18, 25, 26]. We are also working with the Broad Institute to integrate LISA into applications that are widely used by the genomics community.

References

- [1] Illumina, Inc. Novasec 6000 sequencing system. <https://www.illumina.com/content/dam/illumina-marketing/documents/products/datasheets/novaseq-6000-system-specification-sheet-770-2016-025.pdf>, 2019. Accessed: September 2019.
- [2] Vasimuddin Md, Sanchit Misra, and Srinivas Aluru. Identification of Significant Computational Building Blocks through Comprehensive Investigation of NGS Secondary Analysis Methods. *bioRxiv*, April 2018. URL <https://www.biorxiv.org/content/early/2018/07/25/301903>.
- [3] Ben Langmead, Cole Trapnell, Mihai Pop, and Steven L Salzberg. Ultrafast and Memory-efficient Alignment of Short DNA Sequences to the Human Genome. *Genome Biology*, 10(3):1, 2009.
- [4] Ben Langmead and Steven L Salzberg. Fast Gapped-read Alignment with Bowtie 2. *Nature Methods*, 9(4):357–359, 2012.
- [5] Ruiqiang Li, Chang Yu, Yingrui Li, Tak-Wah Lam, Siu-Ming Yiu, Karsten Kristiansen, and Jun Wang. SOAP2: An Improved Ultrafast Tool for Short Read Alignment. *Bioinformatics*, 25(15):1966–1967, 2009.
- [6] Ruibang Luo, Thomas Wong, Jianqiao Zhu, Chi-Man Liu, Xiaoqian Zhu, Edward Wu, Lap-Kei Lee, Haoxiang Lin, Wenjuan Zhu, David W Cheung, et al. SOAP3-dp: Fast, Accurate, and Sensitive GPU-Based Short Read Aligner. *PLOS ONE*, 8(5):e65632, 2013.
- [7] Heng Li and Richard Durbin. Fast and Accurate Short Read Alignment with Burrows–Wheeler Transform. *Bioinformatics*, 25(14):1754–1760, 2009.
- [8] Paolo Ferragina and Giovanni Manzini. Opportunistic Data Structures with Applications. In *IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 390–398, 2000.
- [9] Alejandro Chacón, Juan Carlos Moure, Antonio Espinosa, and Porfidio Hernández. n-step FM-Index for Faster Pattern Matching. *Procedia Computer Science*, 18:70–79, 2013.
- [10] Alejandro Chacón, Santiago Marco-Sola, Antonio Espinosa, Paolo Ribeca, and Juan Carlos Moure. Boosting the FM-Index on the GPU: Effective Techniques to Mitigate Random Memory Access. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 12(5):1048–1059, September 2015.
- [11] J. Pantaleoni and N. Subtil. NVBIO: A Library of Reusable Components Designed by NVIDIA Corporation to Accelerate Bioinformatics Applications using CUDA. <http://nvlabs.github.io/nvbio/>. Accessed: November 2017.
- [12] Jing Zhang, Heshan Lin, Pavan Balaji, and Wu-chun Feng. Optimizing Burrows-Wheeler Transform-based Sequence Alignment on Multi-core Architectures. In *IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid)*, pages 377–384, 2013.
- [13] E. Fernandez, W. Najjar, and S. Lonardi. String Matching in Hardware Using the FM-Index. In *IEEE Annual International Symposium on Field-Programmable Custom Computing Machines*, pages 218–225, May 2011.
- [14] Szymon Grabowski, Marcin Raniszewski, and Sebastian Deorowicz. FM-index for Dummies. In *International Conference on Beyond Databases, Architectures, and Structures (BDAS)*, pages 189–201, 2017.
- [15] Sanchit Misra, Tony C Pan, Kanak Mahadik, George Powley, Priya N. Vaidya, Md Vasimuddin, and Srinivas Aluru. Performance Extraction and Suitability Analysis of Multi- and Many-core Architectures for Next Generation Sequencing Secondary Analysis. In *International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pages 3:1–3:14, 2018.
- [16] Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. The Case for Learned Index Structures. In *ACM International Conference on Management of Data (SIGMOD)*, pages 489–504, June 2018.

- [17] Darryl Ho, Jialin Ding, Sanchit Misra, Nesime Tatbul, Vikram Nathan, Vasimuddin Md, and Tim Kraska. LISA: Towards Learned DNA Sequence Search, 2019. URL <http://arxiv.org/abs/1910.04728>.
- [18] Vasimuddin Md, Sanchit Misra, Heng Li, and Srinivas Aluru. Efficient Architecture-Aware Acceleration of BWA-MEM for Multicore Systems. In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2019.
- [19] Paolo Ferragina and Giovanni Manzini. An Experimental Study of an Opportunistic Index. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 269–278, 2001.
- [20] Alex Galakatos, Michael Markovitch, Carsten Binnig, Rodrigo Fonseca, and Tim Kraska. FITting-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. ACM. ISBN 978-1-4503-5643-5. doi: 10.1145/3299869.3319860. URL <http://doi.acm.org/10.1145/3299869.3319860>.
- [21] Paolo Ferragina and Giorgio Vinciguerra. The PGM-index: a multicriteria, compressed and learned approach to data indexing, 2019. URL <http://arxiv.org/abs/1910.06169>.
- [22] Jialin Ding, Umar Farooq Minhas, Hantian Zhang, Yinan Li, Chi Wang, Badrish Chandramouli, Johannes Gehrke, Donald Kossmann, and David B. Lomet. ALEX: An Updatable Adaptive Learned Index, 2019. URL <http://arxiv.org/abs/1905.08898>.
- [23] Yingjun Wu, Jia Yu, Yuanyuan Tian, Richard Sidle, and Ronald Barber. Designing Succinct Secondary Indexing Mechanism by Exploiting Column Correlations. In *Proceedings of the 2019 International Conference on Management of Data, SIGMOD '19*, pages 1223–1240, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-5643-5. doi: 10.1145/3299869.3319861. URL <http://doi.acm.org/10.1145/3299869.3319861>.
- [24] Michael I. Jordan and Robert A. Jacobs. Hierarchical Mixtures of Experts and the EM Algorithm. *Neural Comput.*, 6(2):181–214, March 1994. ISSN 0899-7667. doi: 10.1162/neco.1994.6.2.181. URL <http://dx.doi.org/10.1162/neco.1994.6.2.181>.
- [25] Heng Li. Exploring Single-Sample SNP and INDEL Calling with Whole-Genome De Novo Assembly. *Bioinformatics*, 28(14):1838–1844, 2012.
- [26] Heng Li. Aligning Sequence Reads, Clone sequences, and Assembly Contigs with BWA-MEM. *arXiv preprint arXiv:1303.3997*, 2013.