

Efficient Tree-Structured Categorical Retrieval

Djamal Belazzougui¹

CAPA, DTISI, Centre de Recherche sur l'Information Scientifique et Technique, Algiers, Algeria
dbelazzougui@cerist.dzm

Gregory Kucherov 

CNRS and LIGM/Univ Gustave Eiffel, Marne-la-Vallée, France
Skolkovo Institute of Science and Technology, Moscow, Russia
Gregory.Kucherov@univ-mlv.fr

Abstract

We study a document retrieval problem in the new framework where D text documents are organized in a *category tree* with a pre-defined number h of categories. This situation occurs e.g. with taxonomic trees in biology or subject classification systems for scientific literature. Given a string pattern p and a category (level in the category tree), we wish to efficiently retrieve the t *categorical units* containing this pattern and belonging to the category. We propose several efficient solutions for this problem. One of them uses $n(\log \sigma(1 + o(1)) + \log D + O(h)) + O(\Delta)$ bits of space and $O(|p| + t)$ query time, where n is the total length of the documents, σ the size of the alphabet used in the documents and Δ is the total number of nodes in the category tree. Another solution uses $n(\log \sigma(1 + o(1)) + O(\log D)) + O(\Delta) + O(D \log n)$ bits of space and $O(|p| + t \log D)$ query time. We finally propose other solutions which are more space-efficient at the expense of a slight increase in query time.

2012 ACM Subject Classification Theory of computation → Pattern matching; Information systems → Document representation; Information systems → Information retrieval query processing; Theory of computation → Data structures design and analysis

Keywords and phrases pattern matching, document retrieval, category tree, space-efficient data structures

Digital Object Identifier 10.4230/LIPIcs.CPM.2020.4

Related Version <http://arxiv.org/abs/2006.01825>

Funding GK was partially funded by RFBR, project 20-07-00652, and joint RFBR and JSPS project 20-51-50007.

1 Introduction

Data is often structured using *category hierarchies* represented by trees. In many applications, such hierarchies play a crucial guiding role: for example, the International Classification of Diseases (ICD) provides a hierarchical classification of all human diseases and constitutes a common reference for diagnostics. In this paper, we are interested in sequence data, such as biological sequences or text documents, that are linked to a given hierarchy. More precisely, in our framework sequences are associated to leaves of a hierarchy, and tree nodes are mapped to several fixed levels, also called ranks.

This situation is common and occurs in several important applications. One is biology where species are classified according to the famous Linnaean taxonomy including eight common *taxonomic ranks*: species, genus, family, order, class, phylum, kingdom, domain. Then, given a set of sequences (DNA, RNA or protein) belonging to known species, one can associate them to the corresponding leaves of the taxonomic tree. Such a structure is used, for example, for phylogeny-based metagenomic classification where one considers the tree of

¹ Corresponding author



known genomic sequences as a reference for classifying sequences of a metagenomic sample, see e.g. [23]. A classification procedure may involve queries asking for the taxonomic units (i.e. internal nodes of the tree) of a certain rank whose sequences contain a given pattern, or similar type of queries.

Another example is provided by text documents such as scientific papers. The latter are usually annotated by subjects belonging to a fixed hierarchical nomenclature, such as ACM Computing Classification System (CCS) or Mathematics Subject Classification (MSC). Those subject hierarchies have a predefined number of levels: four levels for CCS and three for MSC. Given a corpus of scientific papers, one could ask about subject categories at a certain level whose documents contain a given pattern. This is a natural information retrieval scenario.

Here we study this problem from the stringology perspective (see e.g. [14, 8]). Assume we are given a set of D documents of total length n over an alphabet of size σ , organized in a tree of height h . The tree has D leaves, each associated with a distinct document, and the leaves are all at level h of the tree. The total number of nodes in the tree is denoted by Δ . The tree specifies a hierarchy of categories: each level of the tree corresponds to a category, and each internal node corresponds to a *categorical unit*.

The basic type of query we study in this paper is the following.

Given a pattern p , and a tree level (rank) $i \in [1..h]$, return all nodes (categorical units) d_1, \dots, d_t at level i that have at least one leaf (document) in their subtree that contains pattern p .

For example, given a large collection of genomic sequences organized in a taxonomic tree (for example, all known animal genomes), one may ask which animal families have a given sequence in the genomes of their members. Or, given a large hierarchy of documents (for example, all Computer Science papers), one may wonder in which subfields of Computer Science (corresponding to a certain level of the hierarchy) the term 'suffix tree' is used. This basic type of queries can be further extended in different ways. For example, one may impose an additional requirement of the minimum number of documents of the categorical unit containing the given pattern. In this first study, we focus on the basic query type.

In this work, we propose several algorithms for this problem. Our first solution (Section 3) is based on the approach of Muthukrishnan [16] to the document retrieval problem. By combining several algorithmic tools - efficient text index, colored range reporting queries, and level ancestor queries - we obtain a solution with $n(\log \sigma(1 + o(1)) + \log D + O(h)) + O(\Delta)$ bits of space and $O(|p| + t)$ query time, where t is the output size, i.e. the number of retrieved categorical units. To improve the space bound, in particular to get rid of the $O(nh)$ term which can be as big as $O(nD)$, we then develop a solution based on a wavelet tree built on top of the input category tree (Section 4). On this way, we first obtain a solution taking $n(\log \sigma + \log D) + O(D \log n)$ bits and $O(|p| + t \cdot h \log D)$ query time. We further improve it using the technique of heavy path decomposition, to obtain a solution in $n(\log \sigma(1 + o(1)) + \log D) + O(\Delta)$ bits of space and $O(|p| + t \log D)$ query time. In the final part of the paper (Section 5), we focus on solutions using succinct and compressed data structures, on top of the input data. That is, our main goal here is to replace the $n \log D$ bits by respectively $n \log \sigma$ or by $nH_0 + o(n \log \sigma)$ in representing the document array. We obtain memory-time trade-offs showing how this goal can be achieved at the price of a slight increase of query time.

We summarize our main results in the following table.

algorithm	space (bits)	query time
based on colored range queries (Sect. 3)	$n(\log \sigma(1 + o(1)) + \log D + O(h)) + O(\Delta)$	$O(p + t)$
based on wavelet tree (Sect. 4)	$n(\log \sigma(1 + o(1)) + O(\log D)) + O(\Delta) + O(D \log n)$	$O(p + t \log D)$
compact space (Sect. 5)	$O(n \log \sigma)$	$O(p + (t + 1) \cdot \log^\epsilon n(1 + \frac{h}{\log \sigma}))$
compressed space (Sect. 5)	$nH_k + o(n \log \sigma) + O(D \log n)$	$O(p + t \cdot h \log n (\log \log n)^2)$

2 Preliminaries

We first briefly present main algorithmic tools used by our algorithms.

2.1 Level ancestor queries on trees

Consider a rooted tree. To each node in the tree we associate its *level* so that the level of the root is 1, and the level of a child node is 1 more than the level of its parent. The height of a tree is defined as the maximal level of any node in the tree. We denote by ℓ_α the level of a node α .

We will use the implementation of level ancestor queries specified by the following lemma.

► **Lemma 1** ([19]). *There exists a data structure that represents a tree with n nodes within space $2n + o(n)$ and allows answering the following queries in constant time:*

1. *given a level ℓ and a node α at level at least ℓ , return the ancestor node β of α at level ℓ ,*
2. *given an integer i , return the node α where α is the leaf number i in left-to-right order.*

We denote by $\text{LAQ}(\alpha, i)$ the query which asks for the ancestor at level i of node α . We denote by $\text{leafselect}(i)$ the query which returns the i -th leaf of the tree in left to right order.

2.2 rank/select queries and wavelet trees

rank and **select** queries on sequences constitute basic building blocks of many succinct data structures [13]. Given a string $S[1..n]$ on an alphabet Σ , a query $\text{rank}_c(S, i)$, with $c \in \Sigma$ and $i \in [1..n]$, asks for the number of occurrences of c in $S[1..i]$ and $\text{select}_c(S, j)$ asks for the unique position i such that $S[i] = c$ and $\text{rank}_c(S, i) = j$.

Consider first the important case of binary sequences (bitvectors). The following result is well-known, see [18].

► **Lemma 2.** *A bitvector $B[1..n]$ can be represented using $n + o(n)$ bits of space, so that queries **rank** and **select** are answered in constant time.*

In the case of non-binary alphabet, **rank/select** queries can be efficiently answered using *wavelet trees*. The wavelet tree has been formally introduced in [9], but a similar structure has been used earlier [3]. Suppose we are given a sequence S of length n over an alphabet Σ .

The (*binary*) *wavelet tree* is a binary tree representation of S that is defined recursively as follows. Let $\Sigma_0 \neq \emptyset$ and $\Sigma_1 \neq \emptyset$ form a partition of Σ (that is, $\Sigma = \Sigma_0 \cup \Sigma_1$ and $\Sigma_0 \cap \Sigma_1 = \emptyset$). Then the root of the binary wavelet tree will contain a binary vector B , such that $B[i] = 0$ iff $S[i] \in \Sigma_0$. Let the sequence S_0 (resp., S_1) be formed by keeping only the elements of S that belong to Σ_0 (resp., Σ_1), in the same order. Then, the left (resp., right) child is defined

recursively using S_0 (resp., S_1) and a binary partition of Σ_0 (resp., Σ_1). The recursion stops whenever we reach a leaf that corresponds to a singleton subset of Σ . Such nodes will form the leaves of the wavelet tree. We refer the reader to the survey [17] for more details about wavelet trees. We will make use of the following lemma:

► **Lemma 3** ([9]). *The wavelet tree over the alphabet $[1..\sigma]$ can be represented using $n(\log \sigma + o(1)) + O(\sigma \log n)$ bits of space, supporting **rank** and **select** queries in $O(\log \sigma)$ time.*

The definition of binary wavelet tree can be readily generalized to the non-binary case. As in the binary case, to any node α labeled by an interval Σ_α is (implicitly) associated the sequence S_α which is the subsequence of $S[1..n]$ consisting of all characters belonging to Σ_α . If a node α of a wavelet tree has d children, then the alphabet interval $\Sigma_\alpha \subseteq [1..\sigma]$ assigned to α is partitioned into d disjoint subintervals instead of two, and α stores a sequence C_α over alphabet $[1..d]$ of length $|S_\alpha|$ such that $C_\alpha[i] = j$ iff $S_\alpha[j] \in \Sigma_{\alpha_j}$.

2.3 Text indexes

We assume familiarity with main text indexing structures: suffix trees, suffix arrays and BWT-indexes. Here we only recall some basic facts about them.

Given a text T over an alphabet $\Sigma = [1..\sigma]$, a suffix tree [22] is a tree data structure that stores in its leaves the suffixes of $T\$$, where $\$$ is a special character that does not appear in T and is lexicographically smaller than any character of T . Each suffix is associated with its starting position in $T\$$. Suffix tree allows answering basic string pattern matching queries: given a pattern p , return the set of starting positions of p in T .

The suffix array of T is a related but more space-efficient data structure defined as the array $\text{SA}[1..n+1]$ obtained by sorting all the suffixes of $T\$$ in lexicographic order and setting $\text{SA}[i] = j$ if and only if the suffix $T[j..n]\$$ has lexicographic rank i among all suffixes of $T\$$.

A suffix tree occupies $O(n \log n)$ bits of space and a matching query needs access to the original text T in addition to the suffix tree. The query time is $O(|p| \log \sigma)$. The suffix array [15] is an alternative to the suffix tree which occupies the same $O(n \log n)$ bits of space, but has lower constant factors in space and supports matching queries in $O(|p| + \log n)$ time.

The BWT-index (FM-index) is a space-efficient alternative to suffix arrays and suffix trees which uses $O(n \log \sigma)$ bits of space only. It was originally proposed in [4] and has seen many improvements. We will use the following version of BWT-index with alphabet-independent query time.

► **Lemma 4** ([1]). *Given a text T of length n over alphabet $[1..\sigma]$, we can build a BWT-index which occupies $n \log \sigma(1 + o(1))$ bits of space and supports computing the range of suffixes prefixed by a pattern p in time $O(|p|)$.*

Note that computing the range of suffixes answers also whether the pattern occurs in the text at all, and if so, reports the number of its occurrences (the size of the lexicographic order interval). For this reason, the query presented in the lemma above is usually referred to as a **count** query. The BWT-index is usually augmented with position information so that it becomes able to report the location of each occurrence of the pattern in addition to the number of occurrences. This can be achieved using for the example the compressed suffix array representation:

► **Lemma 5** ([10]). *Given a text T of length n over alphabet $[1..\sigma]$ and a constant $\epsilon > 0$, we can build a data structure which occupies $O(n \log \sigma)$ bits of space and that returns $\text{SA}[i]$ for any $i \in [1..n]$ in time $O(\log^\epsilon n)$.*

All the above-mentioned text indexes can trivially be extended to support the same type of queries on a collection of documents instead of a single document. More precisely, given a collection of texts T_1, T_2, \dots, T_D over the same alphabet Σ , the same queries can be supported by constructing an index of the string $T_1\$T_2\$ \dots T_D\$$.

2.4 Colored range reporting and document retrieval

Muthukrishnan [16] was the first to study the problem of efficiently retrieving documents containing a given string pattern. Through the use of a text index, he reduced the problem to the one of *color range reporting*, i.e. reporting all *distinct* values (“colors”) occurring in a given interval of an array. His data structure relies on the use of *range minimum query* data structures – a data structure that can find in constant time the smallest element in an sub-range of an array. His algorithm was subsequently improved in terms of space (Theorem 4 in [20]). We will use the following result on color range reporting, which can be obtained by using the optimal range-minimum query data structure [5] in the method of [20]:

► **Lemma 6.** *Given an array $A[1..n] \in [1..\sigma]^n$, we can build a static data structure that occupies $2n + o(n)$ bits that allows reporting all d distinct values occurring in a query interval $A[i..j]$ in time $O(d)$ ($O(1)$ time per reported value). The query will make read-only access to the data structure, read-only random access to elements of the array A and read-write access to a bitvector B of size σ . The bitvector needs to be initialized to zero before the first query and is reset to zero at the end of each query.*

In combination with text indexing, colored range reporting allows supporting document retrieval queries. More precisely, define the *document array* as follows: given a collection of D documents $T_1, T_2 \dots T_D$ of total length n , lexicographically sort all the suffixes of the text $T^* = T_1\$T_2\$ \dots T_D\$$, and set $A[i] = j$ iff the suffix of T^* of lexicographic rank i starts inside T_j (if the suffix starts with $\$$, then set $A[i] = 0$). Document array A can be easily obtained from a text index of $T^* = T_1\$T_2\$ \dots T_D\$$. For this, one can construct a bitmap of length $|T^*|$ with 1’s at positions of $\$$ in T^* and 0’s otherwise. Then $A[i] = \text{rank}_1(A, \text{SA}[i]) + 1$ for $i > D$ and $A[i] = 0$ for $i \leq D$. It is then immediate that using these data structures, Lemmas 4, 5, and 6 lead to solving the document retrieval problem in time $O(|p| + d \log^\epsilon n)$, where d is the number of resulting documents. For this, we can use the document alphabet-independent BWT index to compute the range $[i..j]$ of occurrences of p in $O(|p|)$ time and then report the d distinct documents that appear in the range $A[i..j]$ in $O(d \log^\epsilon n)$ time.

3 Solution based on Muthukrishnan’s data structure

Our first solution will be a combination of tools presented in the previous section. We first build a text index for the concatenation of documents $T_1\$T_2\$ \dots T_D\$$. More specifically, we build an instance of the text index of Lemma 4 which occupies $n \log \sigma(1 + o(1))$ bits and allows to locate the interval of all suffixes of the documents that start with p in time $O(|p|)$. We also build the document array $A[1..n]$, of size $n \log D$, indexed by the document suffixes sorted in lexicographic order and storing the documents each of the suffixes belongs to.

We further store h instances C_1, \dots, C_h of the data structure of Lemma 6, one instance per level of the tree, defined as follows. Consider d (virtual) arrays $A_i[1..n]$, one per level $i \in [1..h]$ of the tree, such that $A_i[j]$ stores the ancestor at level i of document $A[j]$. Then, each C_i is the data structure of Lemma 6 for supporting range reporting queries on array A_i . Thus, C_i allows to return, for any interval $[r..\ell]$, all distinct elements in $A_i[r..\ell]$ in constant time per element provided that a random-access to each element in A_i is supported in constant time.

Note that according to Lemma 6, a query will need to use D bits of working space² since it will need to use a temporary bitvector B of size $D_i \leq D$ where D_i is the number of nodes at level i of the tree³. By Lemma 6, each C_i occupies only $2n + o(n)$. Finally, in order to simulate constant-time random access to entries of arrays A_i , $1 \leq i \leq h$, we build a data structure for constant-time level ancestor queries on the category tree (Lemma 1). Notice that we can access cell $A_i[j]$ using the formula $A_i[j] = \text{LAQ}(\text{leafselect}(A[j]), i)$. The data structure will occupy $2\Delta + o(\Delta)$ bits of space, where Δ is the total number of nodes in the tree.

To answer a query consisting of a pattern p and level i , we proceed as follows. We first compute, in time $O(|p|)$, the interval $[\ell..r]$ of suffixes using the BWT-index (Lemma 4). The documents containing p are then those contained in $A[\ell..r]$. We then have to output all distinct ancestors at level i of documents $A[\ell..r]$, i.e. all distinct elements of $A_i[\ell..r]$. This is done in constant time per reported element using C_i , as follows from Lemma 6 and constant-time access to elements of A_i using LAQ and leafselect queries.

The document array occupies $n \log D$ bits of space. The text index is built on top of the $n \log \sigma(1 + o(1))$ bits. Each of the h instances of the data structure of Lemma 6 will occupy $2n + o(n)$ bits of space each for a total space of $2nh + o(hn)$ bits of space. The data structure built on top of the category tree occupies $2\Delta + o(\Delta)$ bits of space.

We thus have proved the following theorem:

► **Theorem 7.** *Given a collection of D documents of total length n over alphabet $[1..\sigma]$ so that the documents are organized in a hierarchy of documents represented by a tree of total size Δ and of height h , we can build a data structure of size $n(\log \sigma(1 + o(1)) + \log D + O(h)) + O(\Delta)$ bits of space that, given a pattern p , can find all t categories of documents at a given level i that have at least one document that contains the pattern in total time $O(|p| + t)$.*

This data structure will be good enough whenever h is small, for example, when $h = \log D$, which holds for example when each internal node in the tree has at least two children.

4 Wavelet-tree-based solution

If each node of our tree is branching, i.e. has two or more children, then $h = O(\log D)$ and the solution of Section 3 takes $O(n(\log \sigma + \log D))$ bits of space. (Recall that all leaves of our tree occur at level h) However, this may not be the case as the tree may have many non-branching (unary) nodes. In the extreme case, we may have $h = \Omega(D)$ and the space of Theorem 7 will become $\Omega(nD)$ which can be too large if D is large. In this section, we deal with this issue and present solutions based on wavelet trees.

As in Section 3, we assume that we first located an interval $[\ell..r]$ in the document array A that corresponds to the occurrences of the query pattern p . The goal is then to return all internal nodes at level i containing documents from $A[\ell..r]$ in their subtree. In Section 4.1, we present the first "warm-up" solution that we subsequently improve in Section 4.2.

² We define the working space as a writable space that is only used during queries and is restored to its initial state at the end of the query

³ We can use the same bitvector B (Lemma 6) of size D for all h levels: for a query on level i , the first D_i bits of B are initially set to zero and are reset to zero at the end of the query

4.1 Basic wavelet-tree-based solution

We build our wavelet tree on top of the input tree representing the hierarchy of the documents. Therefore, our initial wavelet tree is generally non-binary and non-balanced. As does the input tree, our wavelet tree has height h and $O(\Delta)$ nodes in total. To save space, we will eliminate unary nodes from the wavelet tree (such a node α stores a trivial sequence $C_\alpha = 1^{|S_\alpha|}$, see Section 2.2) and only encode $O(D)$ branching nodes. For each branching node α we store its depth denoted δ_α . Besides the wavelet tree, we will need a data structure for level ancestor queries (Lemma 1) for the input tree that occupies $O(\Delta)$ bits of space and answers queries in constant time.

Our alphabet Σ will be defined to be the set of documents $[1..D]$. The alphabet interval Σ_α assigned to a node α will be the indices of documents occurring in the subtree rooted at α . The string S for which the tree is built will be the document array $A[1..n]$.

Our wavelet tree may have nodes with more than two children and we implement them by local *binarization*. If a node has d children, we will encode it using a binary wavelet tree of $\log d$ levels, called a *local wavelet tree*. In total, the wavelet tree occupies $n(h \log D)$ bits, since the tree contains h levels and each of the n elements of the document array will contribute at most $\log D$ bits to each level.

Consider now a query which is defined by a pattern p and a level i in the input tree. Once we computed the document array interval corresponding to p , say $A[\ell..r]$, we use our wavelet tree to identify the desired nodes at level i . Starting from the root, we traverse the tree top-down through all the nodes α whose assigned sub-alphabet $\Sigma_\alpha \subseteq [1..D]$ intersects with elements of $A[\ell..r]$. This is done by recomputing the current interval for each traversed node. An invariant of this computation is that querying a node α with an interval $[i..j]$ ensures that all elements of $A[\ell..r] \cap \Sigma_\alpha$ are within $S_\alpha[i..j]$. Interval computation is done using **rank** queries on binary vectors B_α stored at nodes α of the wavelet tree, we refer to [7] where this computation is described in detail. We stop the traversal at a node α as soon as $\delta_\alpha \geq i$ and report its ancestor at level i using the level ancestor data structure.

The original tree has at most h levels and each node is replaced by a local wavelet tree with at most $\log D$ levels, therefore a root-to-leaf path in the wavelet tree has at most $h \log D$ nodes, and the total worst-case query time will be $O(h \log D)$ per reported node.

We now analyse the space usage of the data structure. Since the wavelet tree has D leaves and all nodes are branching, the total number of nodes is $O(D)$. Thus, the total space used by the wavelet trees is $n(h \log D)(1 + o(1)) + O(D \log n)$ bits (see Lemma 3). The space used by the BWT-index is $n \log \sigma(1 + o(1))$ (Lemma 4) and the space used by the document array is $n \log D$ bits. The space used by the data structure for level-ancestor queries is $O(\Delta)$ bits (Lemma 1). We thus proved the following theorem.

► **Theorem 8.** *Given a collection of D documents of total length n over alphabet $[1..\sigma]$ and so that the documents are organized in a hierarchy of documents represented by a tree of height h , we can build a data structure of size $n(\log \sigma + (h + 1) \log D)(1 + o(1)) + O(D \log n) + O(\Delta)$ bits of space that can, given a pattern p , find all t categories of documents at level i that have at least one document that contains the pattern in total time $O(|p| + t \cdot h \log D)$.*

4.2 Solutions based on heavy path decomposition

We now describe a more sophisticated solution based on the *heavy path decomposition* [21, 11] of the wavelet tree from the previous section. Here we present a high-level description of our algorithms, full details will be given in the extended version of the paper.

There are several variants of the definition of heavy path decomposition, with slight differences between the variants. In what follows we will use the following variant. With each node α of a given tree T , we associate a *weight* $w(\alpha)$ equal to the number of leaves in the subtree rooted at α . The *heavy child* β of α is the child of α with the greatest weight, with ties resolved arbitrarily. The other children of α are called *light*. The edge between α and its heavy child is called a *heavy edge*, whereas all the other edges from α to its children are called *light edges*.

The heavy path decomposition of a tree T is a decomposition of T into paths defined recursively as follows. We first compute the heavy path (i.e. a path consisting of heavy edges) from the root of T to a leaf, and then recursively apply the decomposition to all subtrees rooted at all light children of the heavy path nodes. An interesting property of the heavy path decomposition is that the number of light edges on any root-to-leaf path is at most $\log D$, where D is the number of leaves in the tree.

4.2.1 First solution based on heavy path decomposition

Our first solution will be neither space- nor time-optimal. For each heavy path starting at a node α for which the number of light children of nodes of the path is ℓ_α , the alphabet will be of size ℓ_α . We can order the nodes (light children) by increasing depths. The sequence S_α that is associated with a heavy path $\alpha = \alpha_1, \dots, \alpha_k$, will be of length n_α over alphabet $[1..\ell_\alpha]$, where n_α is the total number of occurrences of leaves (documents) in the subtree rooted at α in the document array A . That is, the sequence will be a subsequence of $A[1..n]$, where only the documents that belong to the leaves under α are kept, and the encoding of each element in the subsequence will be the index of the (light) children of the heavy path nodes under which the document appears. Let the depths of the nodes in the heavy path be denoted by $d_1 < \dots < d_k$. We additionally store a bitvector B_α marking the node depths of the different nodes. That is, we initialize the bitvector B_α by all zeros and then set $B_\alpha[d_i] = 1$ for every $i \in [1..k]$.

A query for level i will now proceed as follows. We traverse the tree top-down. For each heavy path, we do the following.

1. We first use the bitvector that marks the node depths to determine a subrange $[1..r]$ of the alphabet that will be used for the query (the light nodes included in the range will have depths at most i , whereas the nodes in the range $[r + 1..h]$ will have depth more than i).
2. We traverse the wavelet tree of the current heavy path. Such a query will spend time $O(t \log \ell_\alpha)$ for a heavy path with ℓ_α light children, in which t distinct light children appear in the sequence.

It is easy to see that the total space will be $O(n \log^2 D)$ bits, since the alphabet size is $O(\log \ell_\alpha)$ for each node α with n_α stored elements and each element of A will incur at most $\log D$ elements in the wavelet trees stored in the heavy paths of the tree. The query time can be bounded to be $O(\log^2 D)$ per reported document by a similar argument (we traverse $\log D$ heavy paths and each traversal costs $\log D$ time).

We thus obtain the following result.

► **Theorem 9.** *Given a collection of D documents of total length n over alphabet $[1..\sigma]$ so that the documents are organized in a hierarchy of documents represented by a tree of total size Δ , we can build a data structure of size $O(n \log^2 D + \Delta)$ bits of space that can, given a pattern p , find all t categories of documents at level i that have at least one document containing the pattern in total time $O(|p| + t \log^2 D)$.*

4.2.2 Second solution based on heavy path decomposition

Our second solution based on heavy path decomposition will rely on a more fine-grained encoding. We will make use of Huffman-shaped wavelet tree [6] for each heavy path, such that the wavelet tree node corresponding to a light node of relative weight w (the weight light node divided by weight of the root of heavy path) will be encoded using $\log(1/w) + O(1)$ bits and the corresponding wavelet tree leaf will be at depth $\log(1/w) + O(1)$. It is now easy to see that the encoding of each element of A will take $O(\log D)$ bits and, furthermore, the cost of a query can be upper-bounded by just $O(\log D)$. Both bounds rely on a telescoping argument. We have the following result.

► **Theorem 10.** *Given a collection of D documents of total length n over alphabet $[1..\sigma]$ and so that the documents are organized in a hierarchy of documents represented by a tree of total size Δ , we can build a data structure of size $n(\log \sigma(1 + o(1)) + O(\log D)) + O(\Delta) + O(D \log n)$ bits of space that can, given a pattern p , find all t categories of documents at a level i that have at least one document that contains the pattern in total time $O(|p| + t \log D)$.*

5 Compact and compressed data structures for categorical data queries

In this section we explore more space-efficient versions of the problem. More in detail, we are interested in studying the problem under succinct and compressed-space constraints. Namely, our aim is to use $O(n \log \sigma)$ bits for the succinct case and $nH_0 + o(n \log \sigma) + O(D \log n)$ bits of space for the compressed case. To achieve this, we will improve the solution of Section 3. More precisely, we avoid the storage of the document array and simulate direct access to the document array using Lemma 5. As a consequence, we can achieve time $O(\log^\epsilon n)$ to get the given document index $A[i]$ for any $i \in [1..n]$. This will reduce the space to represent the document array from $O(n \log D)$ to $O(n \log \sigma)$ bits. Now the space used by the range minimum query data structures will become the bottleneck. To reduce the space usage we will make use of sparsification. More precisely, we will divide the document array into blocks and sample just the values of the A array that are the smallest in each block. The space becomes $O(n/\alpha)$ bits where α is the sparsification factor. For details on how the sparsification is used to simulate the reporting of distinct documents that appear in interval $A[i..j]$, we refer the reader to [2, 12]. Here we just mention that the time per reported document becomes $O(\alpha \log^\epsilon n)$ and entails $O(\alpha)$ accesses to the document array, each of which requires $O(\log^\epsilon n)$ time. We thus have the following result.

► **Theorem 11.** *Given a parameter $\alpha \geq 1$ and a collection of D documents of total length n over alphabet $[1..\sigma]$ and so that the documents are organized in a hierarchy of documents represented by a tree of height h , we can build a data structure of size $O(n \log \sigma) + O(nh/\alpha)$ bits of space that can, given a pattern p , find all t categories of documents at level i that have at least one document that contains the pattern in total time $O(|p| + t \cdot \alpha \log^\epsilon n)$.*

By setting $\alpha = \lceil \frac{h}{\log \sigma} \rceil$ we get space $O(n \log \sigma)$ bits and query time $O(|p| + (t + 1) \log^\epsilon n \cdot (1 + \frac{h}{\log \sigma}))$. We thus have the following corollary.

► **Corollary 12.** *Given a parameter α and collection of D documents of total length n over alphabet $[1..\sigma]$ and so that the documents are organized in a hierarchy of documents represented by a tree of height h , we can build a data structure of size $O(n \log \sigma)$ bits of space that can, given a pattern p , find all t categories of documents at level i that have at least one document that contains the pattern in total time $O(|p| + (t + 1) \cdot \log^\epsilon n(1 + \frac{h}{\log \sigma}))$.*

Whenever $h = \log D$ (e.g. every internal node is branching), the query time simplifies to $O(|p| + (t+1) \cdot \log_{\sigma} D \cdot \log^{\epsilon} n) \in O(|p| + (t+1) \log^{1+\epsilon} n)$. We can also get compressed space. Namely, we can use a compressed suffix array [9] with query time $\log n \log \log n$ and space $nH_k + o(n)$ to represent the document array. We will combine the compressed suffix array with the alphabet-independent variant of BWT-index presented in [1]. We then get an index that uses space $nH_k + o(n \log \sigma)$ with query time $O(|p|)$ to find the suffix array interval of a pattern and $O(\log n \log \log n)$ time to access an element of the suffix array. Notice that we can translate access to a suffix array element to an access to a document array element using $O(D \log n)$ bits of space. Summing up, we get the following theorem.

► **Theorem 13.** *Given a parameter α and a collection of D documents of total length n over alphabet $[1..\sigma]$ and so that the documents are organized in a hierarchy of documents represented by a tree of height h , we can build a data structure of size $nH_k + o(n \log \sigma) + O(D \log n) + O(nh/\alpha)$ bits of space that can, given a pattern p , find all t categories of documents at level i that have at least one document that contains the pattern in total time $O(|p| + t \cdot \alpha \log n \log \log n)$.*

By setting $\alpha = h \cdot \log \log n$, we get space $nH_k + o(n \log \sigma) + O(D \log n)$ bits and query time $O(|p| + t \cdot h \log n (\log \log n)^2)$. The latter becomes $O(|p| + t \log D \log n (\log \log n)^2)$ whenever $h = O(\log D)$.

6 Conclusions

In this paper, we proposed several solutions for the problem of categorical retrieval. Possible extensions of our work include the case when the document hierarchy is a DAG rather than a tree. This situation occurs, for example, with phylogenetic networks. The solution in Section 3 could easily be extended to DAG structured categories if there was an efficient support for level ancestor queries on DAGs. Other possible extensions includes top-k queries in which categories are either ordered by a static order or by the total frequency of the pattern in the documents that belong to the reported categories.

References

- 1 Djamel Belazzougui and Gonzalo Navarro. Alphabet-independent compressed text indexing. *ACM Transactions on Algorithms (TALG)*, 10(4):23, 2014.
- 2 Djamel Belazzougui, Gonzalo Navarro, and Daniel Valenzuela. Improved compressed indexes for full-text document retrieval. *Journal of Discrete Algorithms (JDA)*, 18:3–13, 2013.
- 3 Bernard Chazelle. A functional approach to data structures and its use in multidimensional searching. *SIAM Journal on Computing (SICOMP)*, 17(3):427–462, 1988.
- 4 Paolo Ferragina and Giovanni Manzini. Indexing compressed text. *Journal of the ACM (JACM)*, 52(4):552–581, 2005.
- 5 Johannes Fischer and Volker Heun. Space-efficient preprocessing schemes for range minimum queries on static arrays. *SIAM Journal on Computing (SICOMP)*, 40(2):465–492, 2011.
- 6 Luca Foschini, Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. When indexing equals compression: Experiments with compressing suffix arrays and applications. *ACM Transactions on Algorithms (TALG)*, 2(4):611–639, 2006.
- 7 Travis Gagie, Simon J Puglisi, and Andrew Turpin. Range quantile queries: Another virtue of wavelet trees. In *Proceedings of the 16th International Symposium on String Processing and Information Retrieval (SPIRE)*, pages 1–6. Springer, 2009.
- 8 Pawel Gawrychowski, Gregory Kucherov, Yakov Nekrich, and Tatiana Starikovskaya. Minimal discriminating words problem revisited. In *Proceedings of the 20th International Symposium*

- on *String Processing and Information Retrieval (SPIRE)*, volume 8214 of *Lecture Notes in Computer Science*, pages 129–140. Springer, 2013.
- 9 Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. High-order entropy-compressed text indexes. In *Proceedings of the 14th annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 841–850. Society for Industrial and Applied Mathematics, 2003.
 - 10 Roberto Grossi and Jeffrey Scott Vitter. Compressed suffix arrays and suffix trees with applications to text indexing and string matching. *SIAM Journal on Computing (SICOMP)*, 35(2):378–407, 2005.
 - 11 Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing (SICOMP)*, 13(2):338–355, 1984.
 - 12 Wing-Kai Hon, Rahul Shah, Sharma V Thankachan, and Jeffrey Scott Vitter. Space-efficient frameworks for top-k string retrieval. *Journal of the ACM (JACM)*, 61(2):1–36, 2014.
 - 13 G. Jacobson. Space-efficient static trees and graphs. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS), Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 549–554. IEEE Computer Society, 1989. doi:10.1109/SFCS.1989.63533.
 - 14 Gregory Kucherov, Yakov Nekrich, and Tatiana Starikovskaya. Computing discriminating and generic words. In L. Calderón-Benavides, C.N. González-Caro, E. Chávez, and N. Ziviani, editors, *Proceedings of the 19th International Symposium on String Processing and Information Retrieval (SPIRE)*, volume 7608 of *Lecture Notes in Computer Science*, pages 307–317. Springer Verlag, 2012.
 - 15 Udi Manber and Gene Myers. Suffix arrays: a new method for on-line string searches. *SIAM Journal on Computing (SICOMP)*, 22(5):935–948, 1993.
 - 16 S. Muthu Muthukrishnan. Efficient algorithms for document retrieval problems. In *Proceedings of the 13th annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 657–666. Society for Industrial and Applied Mathematics, 2002.
 - 17 Gonzalo Navarro. Wavelet trees for all. *Journal of Discrete Algorithms (JDA)*, 25:2–20, 2014.
 - 18 Gonzalo Navarro. *Compact data structures: a practical approach*. University of Cambridge, New York, NY, 2016.
 - 19 Gonzalo Navarro and Kunihiro Sadakane. Fully functional static and dynamic succinct trees. *ACM Transactions on Algorithms (TALG)*, 10(3):16, 2014.
 - 20 Kunihiro Sadakane. Succinct data structures for flexible text retrieval systems. *Journal of Discrete Algorithms (JDA)*, 5(1):12–22, 2007.
 - 21 Daniel D Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *Journal of computer and system sciences (JCSS)*, 26(3):362–391, 1983.
 - 22 Peter Weiner. Linear pattern matching algorithms. In *Proceedings of the 14th Annual Symposium on Switching and Automata Theory*, pages 1–11. IEEE, 1973.
 - 23 Derrick E. Wood and Steven L. Salzberg. Kraken: ultrafast metagenomic sequence classification using exact alignments. *Genome Biology*, 15(3):R46, March 2014. doi:10.1186/gb-2014-15-3-r46.