

R-enum: Enumeration of Characteristic Substrings in BWT-runs Bounded Space

Takaaki Nishimoto ✉

RIKEN Center for Advanced Intelligence Project, Tokyo, Japan

Yasuo Tabei ✉

RIKEN Center for Advanced Intelligence Project, Tokyo, Japan

Abstract

Enumerating characteristic substrings (e.g., maximal repeats, minimal unique substrings, and minimal absent words) in a given string has been an important research topic because there are a wide variety of applications in various areas such as string processing and computational biology. Although several enumeration algorithms for characteristic substrings have been proposed, they are not space-efficient in that their space-usage is proportional to the length of an input string. Recently, the *run-length encoded Burrows-Wheeler transform (RLBWT)* has attracted increased attention in string processing, and various algorithms for the RLBWT have been developed. Developing enumeration algorithms for characteristic substrings with the RLBWT, however, remains a challenge. In this paper, we present *r-enum (RLBWT-based enumeration)*, the first enumeration algorithm for characteristic substrings based on RLBWT. R-enum runs in $O(n \log \log(n/r))$ time and with $O(r \log n)$ bits of working space for string length n and number r of runs in RLBWT. Here, r is expected to be significantly smaller than n for highly repetitive strings (i.e., strings with many repetitions). Experiments using a benchmark dataset of highly repetitive strings show that the results of r-enum are more space-efficient than the previous results. In addition, we demonstrate the applicability of r-enum to a huge string by performing experiments on a 300-gigabyte string of 100 human genomes.

2012 ACM Subject Classification Theory of computation → Data compression

Keywords and phrases Enumeration algorithm, Burrows-Wheeler transform, Maximal repeats, Minimal unique substrings, Minimal absent words

Digital Object Identifier 10.4230/LIPIcs.CPM.2021.21

Supplementary Material *Software (Source Code)*: <https://github.com/TNishimoto/renum>
archived at `swh:1:dir:cde748f235b355d80783d13b86eaa048b7c965ea`

Acknowledgements We thank reviewers for their useful comments.

1 Introduction

Enumerating characteristic substrings (e.g., maximal repeats, minimal unique substrings and minimal absent words) in a given string has been an important research topic because there are a wide variety of applications in various areas such as string processing and computational biology. The usefulness of the enumeration of maximal repeats has been demonstrated in lossless data compression [20], bioinformatics [6, 23] and string classification with machine learning models [30, 28]. The enumeration of minimal unique substrings and minimal absent words has shown practical benefits in bioinformatics [24, 1, 14, 16] and data compression [18, 19]. There is therefore a strong need to develop scalable algorithms for enumerating characteristic substrings in a huge string.

The *Burrows-Wheeler transform (BWT)* [13] is for permutation-based lossless data compression of a string, and many enumeration algorithms for characteristic substrings leveraging BWT have been proposed. Okanohara and Tsujii [30] proposed an enumeration algorithm for maximal repeats that uses BWT and an enhanced suffix array [2]. Since their



© Takaaki Nishimoto and Yasuo Tabei;

licensed under Creative Commons License CC-BY 4.0

32nd Annual Symposium on Combinatorial Pattern Matching (CPM 2021).

Editors: Paweł Gawrychowski and Tatiana Starikovskaya; Article No. 21; pp. 21:1–21:21

Leibniz International Proceedings in Informatics



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

algorithm takes linear time to the length n of a string and $O(n \log n)$ working space, applying it to a huge string is computationally demanding. Beller et al. [12] proposed an enumeration algorithm for maximal repeats that uses a *range distinct query* on the BWT of a string in $O(n \log \sigma)$ time and $n \log \sigma + o(n \log \sigma) + O(n)$ bits of working space for alphabet size σ of a string. Since the working space of their algorithm is linearly proportional to the length n of a string, a large amount of space is expected to be consumed for huge strings. Along the same line of research, Belazzougui et al. [9, 7] proposed an algorithm for enumerating characteristic substrings in $O(n \log \sigma)$ time and $n \log \sigma + o(n \log \sigma) + O(\sigma^2 \log^2 n)$ bits of working space, which resulted in the space usage being linearly proportional to the string length. Thus, developing a more space-efficient enumeration algorithm for the characteristic substrings of a string remains a challenging issue.

Run-length BWT (RLBWT) is a recent, popular lossless data compression, and it is defined as a run-length compressed BWT for strings. Thus, the compression performance of RLBWT has been shown to be high, especially for *highly repetitive strings* (i.e., strings with many repetitions) such as genomes, version-controlled documents, and source code repositories. Kempa and Kociumaka [26] showed an upper bound on the size of the RLBWT by using a measure of repetitiveness. Although several compressed data structures and string processing algorithms that use RLBWT have also been proposed (e.g., [8, 21, 29, 4, 25]), no previous algorithms for enumerating characteristic substrings based on RLBWT have been proposed. Such enumeration algorithms are expected to be much more space-efficient than existing algorithms for highly repetitive strings.

Contribution. We present the first enumeration algorithm for characteristic substrings based on RLBWT, which we call *r-enum (RLBWT-based enumeration)*. Following the idea of the previous works [12, 7], r-enum performs an enumeration by simulating traversals of a *Weiner-link tree* (e.g., [9]), which is a trie, each node of which represents a right-maximal repeat in a string T of length n . Each characteristic substring in T corresponds to a node in the Weiner-link tree of T . This is made possible in $O(n \log \log_w(n/r) + occ)$ time and $O(r \log n)$ bits of working space for number r of runs in RLBWT, machine word size $w = \Theta(\log n)$, and number occ of characteristic substrings. For a highly repetitive string such that $r = o(n \log \sigma / \log n)$ holds, r-enum is more space-efficient than the best previous algorithms taking $O(nd)$ time and $|RD| + O(n)$ bits of space, where $|RD|$ is the size of a data structure supporting range distinct queries and computing the LF function in $O(d)$ time; a pair $(|RD|, d)$ can be chosen as $(|RD|, d) = (n \log \sigma + o(n \log \sigma), O(\log \sigma))$ [15] or $(|RD|, d) = (O(n \log \sigma), O(1))$ [9, Lemmas 3.5 and 3.17]. Table 1 summarizes the running time and working space of state-of-the-art algorithms including those by Okanohara and Tsujii (OT method) [30], Beller et al. (BBO method) [12], and Belazzougui and Cunial (BC method) [7] in comparison with our r-enum.

Experiments using a benchmark dataset of highly repetitive strings show that r-enum is more space-efficient than the previous algorithms. In addition, we demonstrate the applicability of r-enum to a huge string by performing experiments on a 300-gigabyte string of 100 human genomes, which has not been shown in the previous work so far.

The outline of this paper is as follows. Section 2 introduces several basic notions, including the Weiner-link tree. In Section 3, we present a traversal algorithm for the Weiner-link tree of T in $O(r \log n)$ bits. Section 4 presents r-enum for finding the corresponding nodes to maximal repeats, minimal unique substrings, and minimal absent words. In Section 5, we slightly modify r-enum such that it outputs each characteristic substring and its occurrences in T instead of the corresponding node to the characteristic substring. Section 6 shows the performance of our method on benchmark datasets of highly repetitive strings.

■ **Table 1** Summary of running time and working space of enumeration algorithms for (i) maximal repeats, (ii) minimal unique substrings (MUSs), and (iii) minimal absent words (MAWs) for each method. Last column represents main data structure used in each algorithm. Input of each algorithm is string T of length n or BWT of T , and each outputted characteristic substring is represented by pointer with $O(\log n)$ bits. We exclude inputs and outputs from working space. In addition, σ is alphabet size of T , $w = \Theta(\log n)$ is machine word size, r is the number of runs in RLBWT of T , and $occ = O(n\sigma)$ [17] is the number of minimal absent words for T . RD means data structure (i) supporting range distinct queries in $O(d)$ time per output element and (ii) computing LF function in $O(d)$ time; $|RD|$ is its size. We can choose $(|RD|, d) = (n \log \sigma + o(n \log \sigma), O(\log \sigma))$ [15] or $(|RD|, d) = (O(n \log \sigma), O(1))$ [9, Lemmas 3.5 and 3.17].

(i) Maximal repeats	Running time	Working space (bits)	Data structures
OT method [30]	$O(n)$	$O(n \log n)$	Enhanced suffix array
[9, Theorem 7.8]	$O(n)$	$O(n \log \sigma)$	BWT and RD
BBO method [12]	$O(nd)$	$ RD + O(n)$	BWT and RD
BC method [7]	$O(nd)$	$ RD + O(\sigma^2 \log^2 n)$	BWT and RD
r-enum (this study)	$O(n \log \log_w(n/r))$	$O(r \log n)$	RLBWT and RD
(ii) MUSs	Running time	Working space (bits)	Data structures
[9]	$O(n)$	$O(n \log \sigma)$	BWT and RD
BC method [7]	$O(nd)$	$ RD + O(\sigma^2 \log^2 n)$	BWT and RD
r-enum (this study)	$O(n \log \log_w(n/r))$	$O(r \log n)$	RLBWT and RD
(iii) MAWs	Running time	Working space (bits)	Data structures
[5]	$O(n + occ)$	$O(n \log n)$	Suffix array
[9, Theorem 7.12]	$O(n + occ)$	$O(n \log \sigma)$	BWT and RD
BC method [7]	$O(nd + occ)$	$ RD + O(\sigma^2 \log^2 n)$	BWT and RD
r-enum (this study)	$O(n \log \log_w(n/r) + occ)$	$O(r \log n)$	RLBWT and RD

2 Preliminaries

Let $\Sigma = \{1, 2, \dots, \sigma\}$ be an ordered alphabet, T be a string of length n over Σ , and $|T|$ be the length of T . Let $T[i]$ be the i -th character of T (i.e., $T = T[1], T[2], \dots, T[n]$), and $T[i..j]$ be the substring of T that begins at position i and ends at position j . For two strings T and P , $T \prec P$ means that T is lexicographically smaller than P . $Occ(T, P)$ denotes all the occurrence positions of P in T , i.e., $Occ(T, P) = \{i \mid i \in [1, n - |P| + 1] \text{ s.t. } P = T[i..(i + |P| - 1)]\}$. We assume that (i) the last character of T is a special character $\$$ not occurring in substring $T[1..n - 1]$, (ii) $|T| \geq 2$, and (iii) every character in Σ occurs at least once in T . For two integers b and e ($b \leq e$), *interval* $[b, e]$ represents the set $\{b, b + 1, \dots, e\}$. Let $\text{substr}(T)$ denote the set of all the distinct substrings of T (i.e., $\text{substr}(T) = \{T[i..j] \mid 1 \leq i \leq j \leq n\}$).

In this paper, characteristic substrings of a string consist of *maximal repeats*, *minimal unique substrings*, and *minimal absent words*. A maximal repeat in T is defined as a substring P satisfying two conditions: (i) it occurs at least twice in the string (i.e., $|Occ(T, P)| \geq 2$), and (ii) either of the left or right extended substrings of it occurs fewer times than it (i.e., $|Occ(T, cP)|, |Occ(T, Pc)| < |Occ(T, P)|$ for any $c \in \Sigma$). A minimal unique substring is defined as substring P satisfying two conditions: (i) it occurs just once (i.e., $|Occ(T, P)| = 1$), and (ii) all the proper substrings of it occur at least twice in the string (i.e., $|Occ(T, P[2..|P|])|, |Occ(T, P[1..|P| - 1])| \geq 2$). A minimal absent word is defined as string P satisfying two conditions: (i) it does not occur in a string (i.e., $|Occ(T, P)| = 0$), and (ii) all the

Sorted circular strings

i	SA	LCP	LF	F	L
1	11	0	2	\$ a b a a b a b a b a	a
2	10	0	8	a \$ a b a a b a b a b	b
3	3	1	9	a b a b a b a \$ a b	a
4	8	1	10	a b a \$ a b a a b a b	b
5	1	3	1	a b a a b a b a b a \$	\$
6	6	3	11	a b a b a \$ a b a a b	b
7	4	5	3	a b a b a b a \$ a b	a
8	9	0	4	b a \$ a b a a b a b	a
9	2	2	5	b a a b a b a b a \$ a	a
10	7	2	6	b a b a \$ a b a a b a	a
11	5	4	7	b a b a b a \$ a b a a	a

■ **Figure 1** Suffix array, LCP array, LF function, BWT, and circular strings for $T = abaabababa\$$.

proper substrings of it occur in the string (i.e., $|Occ(T, P[2..|P])|, |Occ(T, P[1..|P| - 1])| \geq 1$). For convenience, a minimal absent word is sometimes called a substring, although the string is not a substring of T .

Our computation model is a unit-cost word RAM with a machine word size of $w = \Theta(\log_2 n)$ bits. We evaluate the space complexity in terms of the number of machine words. A bitwise evaluation of the space complexity can be obtained with a multiplicative factor of $\log_2 n$. We assume the base-2 logarithm throughout this paper when the base is not indicated.

2.1 Rank and range distinct queries

Let $S \subseteq \{1, 2, \dots, n\}$ be a set of d integers. A *rank* query $\text{rank}(S, i)$ on S returns the number of elements no more than i in S , i.e., $\text{rank}(S, i) = |\{j \mid j \in S \text{ s.t. } j \leq i\}|$. $R_{\text{rank}}(S)$ is a *rank data structure* solving a rank query on S in $O(\log \log_w(n/d))$ time and with $O(dw)$ bits of space [10].

A range distinct query, $\text{RD}(T, b, e)$ on a string T returns a set of 3-tuples (c, p_c, q_c) that consists of (i) a distinct character c in $T[b..e]$, (ii) the first occurrence p_c of the character c for a given interval $[b..e]$ in T , and (iii) the last occurrence q_c of the character c for $[b..e]$ in T . Formally, let $\Sigma(T[b..e])$ be the set of distinct characters in $T[b..e]$, i.e., $\Sigma(T[b..e]) = \{T[i] \mid i \in [b, e]\}$. Then, $\text{RD}(T, b, e) = \{(c, p_c, q_c) \mid c \in \Sigma(T[b..e])\}$, where $p_c = \min(Occ(T, c) \cap [b, e])$, and $q_c = \max(Occ(T, c) \cap [b, e])$. $R_{\text{RD}}(T)$ is a *range distinct data structure* solving a range distinct query on T in $O(|\text{RD}(T, b, e)| + 1)$ time and with $O(n \log \sigma)$ bits of space [11].

2.2 Suffix and longest common prefix arrays

The *suffix array* [27] SA of string T is an integer array of size n such that $\text{SA}[i]$ stores the starting position of the i -th suffix of T in lexicographical order. Formally, SA is a permutation of $\{1, 2, \dots, n\}$ such that $T[\text{SA}[1]..n] \prec \dots \prec T[\text{SA}[n]..n]$. The *longest common prefix array* (LCP array) LCP of T is an integer array of size n such that $\text{LCP}[1] = 0$ and $\text{LCP}[i]$ stores the length of the LCP of the two suffixes $T[\text{SA}[i]..n]$ and $T[\text{SA}[i - 1]..n]$ for $i \in \{2, 3, \dots, n\}$. We call the values in the suffix and LCP arrays *sa-values* and *lcp-values*, respectively. Moreover, let LF be a function such that (i) $\text{SA}[\text{LF}(i)] = \text{SA}[i] - 1$ for any integer $i \in (\{1, 2, \dots, n\} \setminus \{x\})$ and (ii) $\text{SA}[\text{LF}(x)] = n$, where x is an integer such that $\text{SA}[x] = 1$. Figure 1 depicts the suffix array, LCP array, and LF function for a string.

2.3 BWT and RLBWT

The BWT [13] of string T is an array L built by permuting T as follows; (i) all the n rotations of T are sorted in lexicographical order, and (ii) $L[i]$ for any $i \in \{1, 2, \dots, n\}$ is the last character at the i -th rotation in sorted order. Similarly, $F[i]$ for any $i \in \{1, 2, \dots, n\}$ is the first character at the i -th rotation in sorted order. Formally, let $L[i] = T[\text{SA}[\text{LF}(i)]]$ and $F[i] = T[\text{SA}[i]]$.

Since $L[i]$ and $L[j]$ represent two characters $T[\text{SA}[i] - 1]$ and $T[\text{SA}[j] - 1]$ for two distinct integers $i, j \in \{1, 2, \dots, n\}$, the following relation holds between $\text{LF}(i)$ and $\text{LF}(j)$; $\text{LF}(i) < \text{LF}(j)$ if and only if either of two conditions holds: (i) $L[i] \prec L[j]$ or (ii) $L[i] = L[j]$ and $i < j$. Let C be an array of size σ such that $C[c]$ is the number of occurrences of characters lexicographically smaller than $c \in \Sigma$ in string T , i.e., $C[c] = |\{i \mid i \in [1, n] \text{ s.t. } T[i] \prec c\}|$. The following equation holds by the above relation for the LF function on BWT L : $\text{LF}(i) = C[c] + \text{Occ}(L[1..i], L[i])$. We call the equation *LF formula*.

The RLBWT of T is the BWT encoded by run-length encoding, i.e., RLBWT is a partition of L into r substrings $L[\ell(1)..\ell(2) - 1], L[\ell(2)..\ell(3) - 1], \dots, L[\ell(r)..\ell(r+1) - 1]$ such that each substring $L[\ell(i)..\ell(i+1) - 1]$ is a maximal repetition of the same character in L called a *run*. Formally, let $r = 1 + |\{i \mid i \in \{2, 3, \dots, n\} \text{ s.t. } L[i] \neq L[i-1]\}|$, $\ell(1) = 1$, $\ell(r+1) = n + 1$, and $\ell(j) = \min\{i \mid i \in \{\ell(j-1) + 1, \ell(j-1) + 2, \dots, n\} \text{ s.t. } L[i] \neq L[i-1]\}$ for $j \in \{2, 3, \dots, r\}$. Let $\mathcal{S}_{\text{start}}$ denote the set of the starting position of each run in L , i.e., $\mathcal{S}_{\text{start}} = \{\ell(1), \ell(2), \dots, \ell(r)\}$. RLBWT is represented as r pairs $(L[\ell(1)], \ell(1)), (L[\ell(2)], \ell(2)), \dots, (L[\ell(r)], \ell(r))$ using $2rw$ bits. $\sigma \leq r \leq n$ holds since we assume that every character in Σ occurs in T .

$\text{LF}(i) = \text{LF}(i-1) + 1$ holds for an integer $i \in \{2, 3, \dots, n\}$ by LF formula if i is not the starting position of a run in L (i.e., $i \notin \mathcal{S}_{\text{start}}$). Similarly, $\text{LCP}[\text{LF}(i)] = 1 + \text{LCP}[i]$ holds if $i \notin \mathcal{S}_{\text{start}}$.

Figure 1 depicts two arrays L and F for string $T = \text{abaabababa}\$$. Since the BWT of $T = \text{abaabababa}\$$ is $\text{abb}\$\text{baaaaa}$, the RLBWT of the string T is $(a, 1), (b, 2), (\$, 5), (b, 6)$, and $(a, 7)$. The red and blue characters a are adjacent in a run on the BWT L . Hence, $\text{LF}(8) = \text{LF}(7) + 1$ holds by LF formula. Similarly, $\text{LCP}[\text{LF}(8)] = \text{LCP}[8] + 1 = 1$ holds.

2.4 Suffix tree

The *suffix tree* [3] of T is a trie storing all the suffixes of T . Each node v represents the concatenated string on the path from the root of the tree to the node v . Let u_P denote the node representing a substring P of T . The *depth* of node u_P is defined as the length of its string P , i.e., the depth of u_P is $|P|$. Let $\text{children}(P)$ denote the set of strings represented by the children of u_P . Formally, $\text{children}(P) = \{Pc \mid c \in \Sigma \text{ s.t. } Pc \in \text{substr}(T)\}$. We call u_P an *explicit node* if it has at least two children; otherwise, we call u_P an *implicit node*. The root u_ε of the suffix tree is explicit because T contains at least two distinct characters. Let \mathcal{L}_d be the set of substrings represented by all the explicit nodes with depth d . Formally, $\mathcal{L}_d = \{P \mid P \in \text{substr}(T) \text{ s.t. } |P| = d \text{ and } |\text{children}(P)| \geq 2\}$.

The *suffix-tree interval* (a.k.a. *suffix-array interval*) for a substring P of T is an interval $[b, e]$ on the suffix array of T such that $\text{SA}[b..e]$ represents all the occurrence positions of P in string T ; that is, for any integer $p \in \{1, 2, \dots, n\}$, $T[p..p + |P| - 1] = P$ holds if and only if $p \in \{\text{SA}[b], \text{SA}[b+1], \dots, \text{SA}[e]\}$. The suffix-tree interval of the empty string ε is defined as $[1, n]$. Let $\text{interval}(P) = [b, e]$ denote the suffix-tree interval for P .

The *rich representation* [9] $\text{repr}(P)$ for P is a 3-tuple $(\text{interval}(P), \{(c_1, b_1, e_1), (c_2, b_2, e_2), \dots, (c_k, b_k, e_k)\}, |P|)$. Here, Pc_1, Pc_2, \dots, Pc_k are strings represented by the children of node u_P , and $[b_i, e_i] = \text{interval}(Pc_i)$ for $i \in [1, k]$.

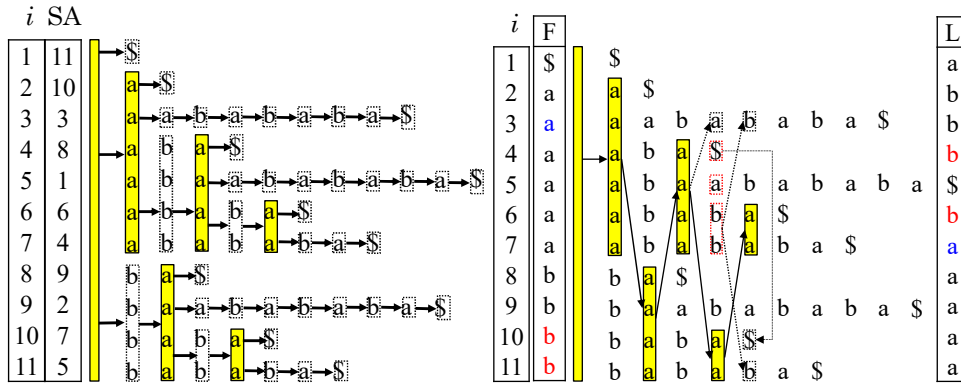


Figure 2 Left figure illustrates suffix tree of $T = abaabababa\$$ on sorted suffixes of T . We represent each node by its suffix-tree interval (rectangles) and omit characters on edges. Yellow and white rectangles are explicit and implicit nodes, respectively. Right figure illustrates Weiner-link tree of T and Weiner links on sorted suffixes of T . Tree consists of yellow rectangles and solid arrows. We omit characters on Weiner links. Solid and dotted arrows represent Weiner links pointing to explicit and implicit nodes, respectively. Red rectangles are children of node of aba in suffix tree of T .

Figure 2 illustrates the suffix tree of string $T = abaabababa\$$ (left figure). The three sets $\mathcal{L}_0, \mathcal{L}_1, \mathcal{L}_2$ are $\{\varepsilon\}, \{a\}, \{ba\}$, respectively. The suffix-tree interval for substring $P = aba$ is $[4, 7]$, and $\text{children}(P) = \{aba\$, abaa, abab\}$. The rich representation for P is $([4, 7], \{(\$, 4, 4), (a, 5, 5), (b, 6, 7)\}, 3)$.

2.5 Weiner links and Weiner-link tree

Weiner links are directed links on the suffix tree of T . Let cP be a substring of T , where c is a character, and P is a string. Then, node u_{cP} is the destination of a Weiner link with character c starting at node u_P . Hence, every node u_P must be the destination of exactly one Weiner link unless node u_P is the root of the suffix tree (i.e., $P = \varepsilon$). Node u_P is always explicit if u_{cP} is explicit. This is because the explicit node u_{cP} indicates that T has two substrings cPc_1 and cPc_2 , where c_1 and c_2 are distinct characters. Let $\text{WLink}(P)$ denote the set of strings such that the node for each string is the destination of a Weiner link starting at node u_P (i.e., $\text{WLink}(P) = \{cP \mid c \in \Sigma \text{ s.t. } cP \in \text{substr}(T)\}$).

The Weiner-link tree (a.k.a. *suffix-link tree*) for T is a graph such that (i) the nodes are all the explicit nodes in the suffix tree of T and (ii) the edges are all the Weiner links among the explicit nodes. Since any explicit node is the destination of a Weiner link starting at another explicit node, the graph results in a tree. Each child of a node u_P represents a string in $\text{WLink}(P)$ in the Weiner-link tree.

Let p_c and q_c be the first and last occurrences of a character c on the suffix-tree interval for a substring P in the BWT L of T , respectively. Then, $[\text{LF}(p_c), \text{LF}(q_c)]$ is equal to the suffix-tree interval of cP . Hence, we can compute the suffix-tree intervals for the destinations of all the Weiner links starting at node u_P using a range distinct query and LF function. Formally, let $[b, e] = \text{interval}(P)$ for a substring P of T . Then, the following equation holds:

$$\{\text{interval}(cP) \mid cP \in \text{WLink}(P)\} = \{[\text{LF}(p_c), \text{LF}(q_c)] \mid (c, p_c, q_c) \in \text{RD}(L, b, e) \text{ s.t. } c \neq \$\} \quad (1)$$

Let Q_P denote an array of size σ for a substring P of T such that $Q_P[c]$ stores set $\{(c', b, e) \mid cPc' \in \text{children}(cP) \text{ and let } [b, e] = \text{interval}(cPc')\}$ for a character $c \in \Sigma$ if cP is a substring of T ; otherwise, $Q_P[c] = \emptyset$. The array Q_P has three properties for any

character $c \in \Sigma$: (i) $cP \in \text{WLink}(P)$ holds if and only if $|Q_P[c]| \geq 1$ holds, (ii) node u_{cP} is explicit if and only if $|Q_P[c]| \geq 2$ holds, and (iii) $Q_P[c] = \bigcup_{Pc' \in \text{children}(P)} \{(c', b, e) \mid \hat{c}Pc' \in \text{WLink}(Pc') \text{ s.t. } \hat{c} = c \text{ and let } [b, e] = \text{interval}(cPc')\}$ holds. In other words, each child of node u_{cP} is the destination of a Weiner link starting at a child of node u_P in the suffix tree.

We can compute the children of a node u_P in the Weiner-link tree using Q_P without explicitly constructing the Weiner-link tree. Formally, the following lemma holds.

► **Lemma 1** (Lemma 4.1 in [9]). *Let RD be a data structure supporting a range distinct query on BWT L in $O((1+k)t_{\text{RD}})$ time and computing LF function in $O(t_{\text{LF}})$ time. Here, k is the number of elements output by the range distinct query. We can compute set $\mathcal{Y} = \{\text{repr}(cP) \mid cP \in \text{WLink}(P)\}$ using (i) $\text{repr}(P)$, (ii) data structure RD, and (iii) an empty array X of size σ . After that, we can divide the set \mathcal{Y} into two sets for explicit and implicit nodes. The computation time and working space are $O(h(t_{\text{RD}} + t_{\text{LF}}))$ and $O((\sigma + h')w)$ bits, respectively, where $h = \sum_{Pc' \in \text{children}(P)} |\text{WLink}(Pc')|$, and $h' = \sum_{cP \in \text{WLink}(P) \cap \mathcal{L}_{|P|+1}} |\text{children}(cP)|$.*

Proof. We compute the outputs with the following three steps. At the first step, we compute set $\{\text{interval}(cP) \mid cP \in \text{WLink}(P)\}$ and convert the empty array X into Q_P using Equation 1 and the third property of Q_P . At the second step, we output the rich representation $\text{repr}(cP) = (\text{interval}(cP), Q_P[c], |P| + 1)$ for each $cP \in \text{WLink}(P)$ and divide the rich representations into two sets for explicit and implicit nodes using the second property of Q_P . At the last step, we remove all the elements from $Q_P[c]$ for each $cP \in \text{WLink}(P)$ to recover X from Q_P . We perform the three steps using range distinct queries to the suffix-tree intervals for node u_P and its children, and the intervals stored in $\text{repr}(P)$. Hence, the running time is $O(h(t_{\text{RD}} + t_{\text{LF}}))$ in total.

Next, we analyze the working space. The rich representation for a node $u_{P'}$ takes $O(w)$ bits if $u_{P'}$ is implicit because it has at most one child. Otherwise, $\text{repr}(P')$ takes $O(|\text{children}(P')|w)$ bits. Hence, the working space is $O((\sigma + h')w)$ bits in total. ◀

Figure 2 illustrates the Weiner-link tree of string $T = \text{abaabababa}\$$ (right figure). Since T contains two substrings $aaba$ and $baba$, the node of aba has two Weiner links pointing to the nodes of $aaba$ and $baba$, i.e., $\text{WLink}(aba) = \{aaba, baba\}$. The Weiner-link tree contains the node of $baba$ but not that of $aaba$ because the former and latter nodes are explicit and implicit, respectively.

The three suffix-tree intervals for aba , $aaba$, and $baba$ are $[4, 7]$, $[3, 3]$, and $[10, 11]$, respectively. A range distinct query on $\text{interval}(aba)$ in the BWT of T returns the set $\{(\$, 5, 5), (a, 7, 7), (b, 4, 6)\}$. Hence, $\text{interval}(aaba) = [\text{LF}(7), \text{LF}(7)]$ and $\text{interval}(baba) = [\text{LF}(4), \text{LF}(6)]$ hold by Equation 1 (See also red and blue characters on two arrays F and L in Figure 2). Figure 2 also illustrates the children of the node of aba in the suffix tree and Weiner links starting from the children. In this example, $Q_P[\$] = \emptyset$, $Q_P[a] = \{(b, 3, 3)\}$, and $Q_P[b] = \{(\$, 10, 10), (b, 11, 11)\}$, where $P = aba$. Hence, the node of $aaba$ has one child, and the node of $baba$ has two children in the suffix tree.

3 Traversing Weiner-link tree in $O(rw)$ bits of space

In this section, we present a breadth-first traversal algorithm for the Weiner-link tree of T in $O(rw)$ bits of working space. The traversal algorithm outputs n sets $\{\text{repr}(P) \mid P \in \mathcal{L}_0\}$, $\{\text{repr}(P) \mid P \in \mathcal{L}_1\}$, \dots , $\{\text{repr}(P) \mid P \in \mathcal{L}_{n-1}\}$ in left-to-right order. Here, each set $\{\text{repr}(P) \mid P \in \mathcal{L}_t\}$ represents the set of the rich representations for all the nodes with depth t in the Weiner-link tree. Hence each node u_P is represented as its rich representation $\text{repr}(P)$.

3.1 Data structures

Let D_{LF} be an array of size r such that $D_{\text{LF}}[i] = \text{LF}(\ell(i))$ for $i \in \{1, 2, \dots, r\}$ and $L' = L[\ell(1)], L[\ell(2)], \dots, L[\ell(r)]$. Our traversal algorithm uses six data structures: (i) the RLBWT of T (i.e., r pairs $(L[\ell(1)], \ell(1)), (L[\ell(2)], \ell(2)), \dots, (L[\ell(r)], \ell(r))$, which are introduced in Section 2.3), (ii) the rich representation for the root of the Weiner-link tree of T (i.e., $\text{repr}(\varepsilon)$), (iii) D_{LF} , (iv) $R_{\text{rank}}(\mathcal{S}_{\text{start}})$, (v) $R_{\text{RD}}(L')$, and (vi) an empty array X of size σ , where $R_{\text{rank}}(\mathcal{S}_{\text{start}})$ is the rank data structure introduced in Section 2.1, and it is built on the set $\mathcal{S}_{\text{start}} = \{\ell(1), \ell(2), \dots, \ell(r)\}$ introduced in Section 2.3. Similarly, $R_{\text{RD}}(L')$ is the range distinct data structure introduced in Section 2.1, and it is built on the string L' . The six data structures require $O((r + \sigma)w)$ bits in total. We construct the six data structures in $O(n \log \log_w(n/r))$ time and $O(rw)$ bits of working space by processing the RLBWT of T (see Appendix A).

We use the two data structures D_{LF} and $R_{\text{rank}}(\mathcal{S}_{\text{start}})$ to compute LF function. Let x be the index of a run containing the i -th character of BWT L (i.e., $x = \text{rank}(\mathcal{S}_{\text{start}}, i)$) for an integer $i \in \{1, 2, \dots, n\}$. $\text{LF}(i) = \text{LF}(\ell(x)) + |\text{Occ}(L[\ell(x)..i], L[i])| - 1$ holds by LF formula, and $|\text{Occ}(T[\ell(x)..i], L[i])| = i - \ell(x) + 1$ holds because $L[\ell(x)..i]$ consists of a repetition of the i -th character $L[i]$. Hence, $\text{LF}(i) = D_{\text{LF}}[x] + (i - \ell(x))$ holds, and we can compute LF function in $O(\log \log_w(n/r))$ time using D_{LF} and $R_{\text{rank}}(\mathcal{S}_{\text{start}})$.

We use the fifth data structure $R_{\text{RD}}(L')$ to compute a range distinct query on the BWT L of T . Let b' and e' be the indexes of the two runs on L containing two characters $L[b]$ and $L[e]$, respectively (i.e., $b' = \text{rank}(\mathcal{S}_{\text{start}}, b)$ and $e' = \text{rank}(\mathcal{S}_{\text{start}}, e)$), for an interval $[b, e] \subseteq \{1, 2, \dots, n\}$. Then the following relation holds between two sets $\text{RD}(L, b, e)$ and $\text{RD}(L', b', e')$.

► **Lemma 2.** $\text{RD}(L, b, e) = \{(c, \max\{\ell(p), b\}, \min\{\ell(q+1) - 1, e\}) \mid (c, p, q) \in \text{RD}(L', b', e')\}$ holds.

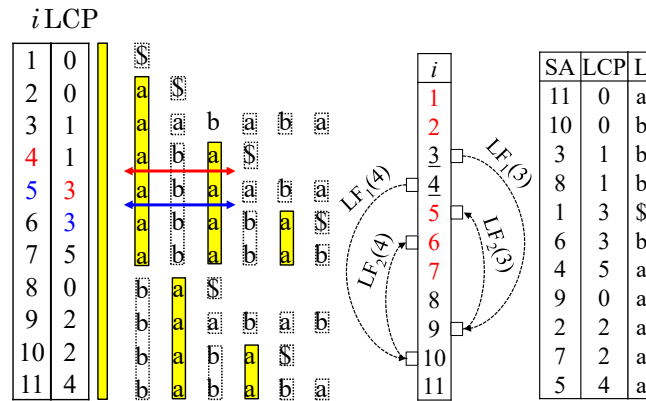
Proof. $L[b..e]$ consists of $e' - b' + 1$ repetitions $L[b..\ell(b' + 1) - 1]$, $L[\ell(b' + 1).. \ell(b' + 2) - 1]$, \dots , $L[\ell(e' - 1).. \ell(e') - 1]$, $L[\ell(e')..e]$. Hence, the following three statements hold: (i) The query $\text{RD}(L', b', e')$ returns all the distinct characters in $L[b..e]$ (i.e., $\{L[i] \mid i \in [b, e]\} = \{L'[i] \mid i \in [b', e']\}$). (ii) Let p be the first occurrence of a character c on $[b', e']$ in L' . Then, the first occurrence of a character c on $[b, e]$ in L is equal to $\ell(p)$ if $\ell(p) \geq b$; otherwise, the first occurrence is equal to b . (iii) Similarly, let q be the last occurrence of the character c on $[b', e']$ in L' . Then, the last occurrence of c on $[b, e]$ in L is equal to $\ell(q + 1) - 1$ if $\ell(q + 1) - 1 \leq e$; otherwise, the last occurrence is equal to e . We obtain Lemma 2 with the three statements. ◀

Lemma 2 indicates that we can solve a range distinct query on L using two rank queries on $\mathcal{S}_{\text{start}}$ and a range distinct query on L' . The range distinct query on L takes $O(\log \log_w(n/r) + k)$ time, where k is the number of output elements.

We use the empty array X to compute rich representations using Lemma 1. The algorithm of Lemma 1 takes $O((k + 1) \log \log_w(n/r))$ time using the six data structures since we can compute LF function and solve the range distinct query on L in $O(\log \log_w(n/r))$ and $O((k + 1) \log \log_w(n/r))$ time, respectively.

3.2 Algorithm

The basic idea behind our breadth-first traversal is to traverse the Weiner-link tree without explicitly building the tree in order to reduce the working space. The algorithm computes nodes sequentially using Lemma 1. Recall that the algorithm of Lemma 1 returns the rich



■ **Figure 3** Left figure illustrates part of suffix tree of $T = abaabababa\$$ on sorted suffixes of T . Two colored integers in array i are positions in $\mathcal{K}'_2 = \{4, 5\}$. Similarly, two colored integers in LCP array correspond to positions in $\mathcal{P}_2 = \{5, 6\}$. Right figure illustrates array i , SA, LCP array, and BWT L for $T = abaabababa\$$. Red integers in array i are integers in set $\mathcal{S}_{\text{start}} = \{1, 2, 5, 6, 7\}$. Similarly, underlined integers are integers in set $\mathcal{P}_1 = \{3, 4\}$.

representations for all the children of a given explicit node u_P in the Weiner-link tree (i.e., it returns set $\{\text{repr}(cP) \mid cP \in \text{WLink}(P) \text{ s.t. } cP \in \mathcal{L}_{|P|+1}\}$). This fact indicates that the set of the rich representations for all the nodes with a depth $t \geq 1$ is equal to the union of the sets of rich representations obtained by applying Lemma 1 to all the nodes with depth $t - 1$ in the Weiner-link tree, i.e., $\{\text{repr}(P) \mid P \in \mathcal{L}_t\} = \bigcup_{P \in \mathcal{L}_{t-1}} \{\text{repr}(cP) \mid cP \in \text{WLink}(P)\}$.

Our algorithm consists of $(n-1)$ steps. At each t -th step, the algorithm (i) applies Lemma 1 to each representation in set $\{\text{repr}(P) \mid P \in \mathcal{L}_{t-1}\}$, (ii) outputs set $\{\text{repr}(P) \mid P \in \mathcal{L}_t\}$, and (iii) removes the previous set $\{\text{repr}(P) \mid P \in \mathcal{L}_{t-1}\}$ from working memory. The algorithm can traverse the whole Weiner-link tree in breadth-first order because we initially store the first set $\{\text{repr}(P) \mid P \in \mathcal{L}_0\}$ for the first step.

3.3 Analysis

The traversal algorithm requires $O(H_{t-1} \log \log_w(n/r))$ computation time at the t -th step. Here, H_t is the number of Weiner links starting from the children of the explicit nodes with depth t (i.e., $H_t = \sum_{P \in \mathcal{L}_t} \sum_{Pc \in \text{children}(P)} |\text{WLink}(Pc)|$). The running time is $O((\sum_{t=0}^n H_t) \log \log_w(n/r))$ in total. The term $\sum_{t=0}^n H_t$ represents the number of Weiner links starting from the children of all the explicit nodes in the suffix tree. Belazzougui et al. showed that $\sum_{t=0}^n H_t = O(n)$ holds [9, Observation 1]. Hence our traversal algorithm runs in $O(n \log \log_w(n/r))$ time.

Next, we analyze the working space of the traversal algorithm. Let \mathcal{K}_t be the set of the children of all the explicit nodes with depth t in the suffix tree of T (i.e., $\mathcal{K}_t = \bigcup_{P \in \mathcal{L}_t} \text{children}(P)$). Then, the algorithm requires $O((r + \sigma + |\mathcal{K}_{t-1}| + |\mathcal{K}_t|)w)$ bits of working space while executing the t -th step. Hence our traversal algorithm requires $O((r + \sigma + \max\{|\mathcal{K}_0|, |\mathcal{K}_1|, \dots, |\mathcal{K}_n|\})w)$ bits of working space while the algorithm runs.

We introduce a set \mathcal{K}'_t to analyze the term $|\mathcal{K}_t|$ for an integer t . The set \mathcal{K}'_t consists of the children of all the explicit nodes with depth t except for the last child of each explicit node, i.e., $\mathcal{K}'_t = \bigcup_{P \in \mathcal{L}_t} \{Pc \mid Pc \in \text{children}(P) \text{ s.t. } e' \neq e\}$, where $[b, e] = \text{interval}(P)$, and $[b', e'] = \text{interval}(Pc)$. $|\mathcal{K}_t| = |\mathcal{K}'_t| + |\mathcal{L}_t|$ holds because every explicit node has exactly one last child. $|\mathcal{L}_t| \leq |\mathcal{K}'_t|$ also holds because every explicit node has at least two children in the suffix tree of T . Hence, we obtain the inequality $|\mathcal{K}_t| \leq 2|\mathcal{K}'_t|$.

We also introduce a set \mathcal{P}_t for an integer t . The set \mathcal{P}_t consists of positions with lcp-value t on the LCP array of T except for the first lcp-value $LCP[1]$, i.e., $\mathcal{P}_t = \{i \mid i \in \{2, 3, \dots, n\} \text{ s.t. } LCP[i] = t\}$. Let P be the longest common prefix of $T[SA[i-1]..n]$ and $T[SA[i]..n]$ for position $i \in \mathcal{P}_t$, i.e., $P = T[SA[i]..SA[i] + LCP[i] - 1]$. Then node u_P is explicit, and two nodes u_{P_c} and $u_{P_{c'}}$ are adjacent children of u_P , where $c = T[SA[i-1] + LCP[i]]$ and $c' = T[SA[i] + LCP[i]]$. The two adjacent children u_{P_c} and $u_{P_{c'}}$ ($P_c \prec P_{c'}$) of an explicit node u_P indicate that $LCP[i] = |P|$ holds, where i is the left boundary of $\text{interval}(P_{c'})$. Hence, there exists a one-to-one correspondence between \mathcal{P}_t and \mathcal{K}'_t , and $|\mathcal{P}_t| = |\mathcal{K}'_t|$ and $|\mathcal{K}_t| \leq 2|\mathcal{K}'_t|$. We obtain the inequality $|\mathcal{K}_t| \leq 2|\mathcal{P}_t|$ by $|\mathcal{P}_t| = |\mathcal{K}'_t|$ and $|\mathcal{K}_t| \leq 2|\mathcal{K}'_t|$.

In Figure 3, the left figure represents a part of a suffix tree. In this example, $\mathcal{P}_3 = \{5, 6\}$, and $\mathcal{K}'_3 = \{4, 5\}$. Obviously, the two positions 5 and 6 in \mathcal{P}_3 correspond to the two positions 4 and 5 in \mathcal{K}'_3 with the adjacent children of explicit node u_{aba} , respectively.

Next, we introduce two functions LF_x and $\text{dist}(i)$ to analyze $|\mathcal{P}_t|$. The first function $\text{LF}_x(i)$ returns the position obtained by recursively applying LF function to i x times, i.e., $\text{LF}_0(i) = i$ and $\text{LF}_x(i) = \text{LF}_{x-1}(\text{LF}(i))$ for $x \geq 1$. The second function $\text{dist}(i) \geq 0$ returns the smallest integer such that $\text{LF}_{\text{dist}(i)}(i)$ is the starting position of a run on L (i.e., $\text{LF}_{\text{dist}(i)}(i) \in \mathcal{S}_{\text{start}}$). Formally, $\text{dist}(i) = \min\{x \mid x \geq 0 \text{ s.t. } \text{LF}_x(i) \in \mathcal{S}_{\text{start}}\}$. The following lemma holds.

► **Lemma 3.** $\text{LF}_{\text{dist}(i)}(i) \neq \text{LF}_{\text{dist}(j)}(j)$ holds for two distinct integers $i, j \in \mathcal{P}_t$, where $t \geq 0$ is an integer.

Proof. We show that $LCP[i] = LCP[\text{LF}_{\text{dist}(i)}(i)] - \text{dist}(i)$ holds for an integer $i \in \{1, 2, \dots, n\}$. Let t be an integer in $[0, \text{dist}(i) - 1]$. Since $\text{LF}_t(i) \notin \mathcal{S}_{\text{start}}$, $LCP[\text{LF}_t(i)] = LCP[\text{LF}_{t+1}(i)] - 1$ holds by LF formula. The LF formula produces $\text{dist}(i)$ equations $LCP[\text{LF}_0(i)] = LCP[\text{LF}_1(i)] - 1$, $LCP[\text{LF}_1(i)] = LCP[\text{LF}_2(i)] - 1$, \dots , $LCP[\text{LF}_{\text{dist}(i)-1}(i)] = LCP[\text{LF}_{\text{dist}(i)}(i)] - 1$. Hence, $LCP[i] = LCP[\text{LF}_{\text{dist}(i)}(i)] - \text{dist}(i)$ holds by the $\text{dist}(i)$ equations.

Next, we prove Lemma 3. The two integers i and j must be the same if $\text{LF}_{\text{dist}(i)}(i) = \text{LF}_{\text{dist}(j)}(j)$ because $\text{dist}(i) = \text{dist}(j)$ holds by three equations $LCP[i] = LCP[j]$, $LCP[i] = LCP[\text{LF}_{\text{dist}(i)}(i)] - \text{dist}(i)$, and $LCP[j] = LCP[\text{LF}_{\text{dist}(j)}(j)] - \text{dist}(j)$. The equation $i = j$ contradicts the fact that $i \neq j$. Hence, $\text{LF}_{\text{dist}(i)}(i) \neq \text{LF}_{\text{dist}(j)}(j)$ holds. ◀

The function $\text{LF}_{\text{dist}(i)}(i)$ maps the integers in \mathcal{P}_t into distinct integers in set $\mathcal{S}_{\text{start}}$ by Lemma 3. The mapping indicates that $|\mathcal{P}_t| \leq |\mathcal{S}_{\text{start}}| = r$ holds for any integer t . In Figure 3, the right figure represents the mapping between $\mathcal{P}_1 = \{3, 4\}$ and $\mathcal{S}_{\text{start}} = \{1, 2, 5, 6, 7\}$ on a BWT. In this example, $\text{LF}_1(3) = 9 \notin \mathcal{S}_{\text{start}}$, $\text{LF}_2(3) = 5 \in \mathcal{S}_{\text{start}}$, $\text{LF}_1(4) = 10 \notin \mathcal{S}_{\text{start}}$, and $\text{LF}_2(4) = 6 \in \mathcal{S}_{\text{start}}$. Hence, $\text{LF}_{\text{dist}(i)}(i)$ maps the two positions 3 and 4 in \mathcal{P}_1 into the two positions 5 and 6 in $\mathcal{S}_{\text{start}}$, respectively, which indicates that $|\mathcal{P}_1| \leq |\mathcal{S}_{\text{start}}|$ holds.

Finally, we obtain $\max\{|\mathcal{K}_0|, |\mathcal{K}_1|, \dots, |\mathcal{K}_n|\} \leq 2r$ by $|\mathcal{K}_t| \leq 2|\mathcal{P}_t|$ and $|\mathcal{P}_t| \leq r$. Hence, the working space of our traversal algorithm is $O((r + \sigma)w)$ bits. We obtain the following theorem using $\sigma \leq r$.

► **Theorem 4.** We can output n sets $\{\text{repr}(P) \mid P \in \mathcal{L}_0\}$, $\{\text{repr}(P) \mid P \in \mathcal{L}_1\}$, \dots , $\{\text{repr}(P) \mid P \in \mathcal{L}_{n-1}\}$ in left-to-right order in $O(n \log \log_w(n/r))$ time and $O(rw)$ bits of working space by processing the RLBWT of T .

4 Enumeration of characteristic substrings in $O(rw)$ bits of space

In this section, we present r-enum, which enumerates maximal repeats, minimal unique substrings, and minimal absent words using RLBWT. While the enumeration algorithm proposed by Belazzougui and Cunial [7] finds nodes corresponding to characteristic substrings by traversing the Weiner-link tree of T in a depth-first manner, r-enum adopts the breadth-first search presented in Section 3. The next theorem holds by assuming $\sigma \leq r$.

► **Theorem 5.** *R-enum can enumerate (i) maximal repeats, (ii) minimal unique substrings, and (iii) minimal absent words for T in $O(n \log \log_w(n/r))$, $O(n \log \log_w(n/r))$, and $O(n \log \log_w(n/r) + occ)$ time, respectively, by processing the RLBWT of T with $O(rw)$ bits of working space, where occ is the number of minimal absent words for T , and $occ = O(\sigma n)$ holds [17]. Here, r -enum outputs rich representation $\text{repr}(P)$, pair $(\text{interval}(P'), |P'|)$, and 3-tuple $(\text{interval}(P''), |P''|, c)$ for a maximal repeat P , minimal unique substring P' , and minimal absent word $P''c$, respectively, where P, P', P'' are substrings of T , and c is a character.*

We prove Theorem 5(i) for maximal repeats. R-enum leverages the following relation between the maximal repeats and nodes in the Weiner-link tree.

► **Lemma 6.** *A substring P of T is a maximal repeat if and only if P satisfies two conditions: (i) node u_P is explicit (i.e., $P \in \mathcal{L}_{|P|}$), and (ii) $\text{rank}(\mathcal{S}_{\text{start}}, b) \neq \text{rank}(\mathcal{S}_{\text{start}}, e)$, where $\text{interval}(P) = [b, e]$.*

Proof. $|\text{Occ}(T, P)| \geq 2$ and $|\text{Occ}(T, P)| > |\text{Occ}(T, Pc)|$ hold for any character $c \in \Sigma$ if and only if node u_P is explicit. From the definition of BWT, $|\text{Occ}(T, P)| > |\text{Occ}(T, cP)|$ holds for any character $c \in \Sigma$ if and only if $L[b..e]$ contains at least two distinct characters, i.e., $\text{rank}(\mathcal{S}_{\text{start}}, b) \neq \text{rank}(\mathcal{S}_{\text{start}}, e)$ holds. Hence, Lemma 6 holds. ◀

R-enum traverses the Weiner-link tree of T , and it verifies whether each explicit node u_P represents a maximal repeat using Lemma 6, i.e., it verifies $\text{rank}(\mathcal{S}_{\text{start}}, b) \neq \text{rank}(\mathcal{S}_{\text{start}}, e)$ or $\text{rank}(\mathcal{S}_{\text{start}}, b) = \text{rank}(\mathcal{S}_{\text{start}}, e)$ using the two rank queries on $\mathcal{S}_{\text{start}}$. We output its rich representation $\text{repr}(P)$ if P is a maximal repeat. The rich representation $\text{repr}(P)$ for u_P stores $[b, e]$, and our breadth-first traversal algorithm stores the data structure $R_{\text{rank}}(\mathcal{S}_{\text{start}})$ for rank queries on $\mathcal{S}_{\text{start}}$. Hence, we obtain Theorem 5(i).

Similarly, r -enum can also find the nodes corresponding to minimal unique substrings and minimal absent words using their properties while traversing the Weiner-link tree. See Appendixes B and C for the proofs of Theorem 5(ii) and (iii), respectively.

5 Modified enumeration algorithm for original characteristic substrings and their occurrences

Let $\text{output}(P)$ be the element representing a characteristic substring P outputted by r -enum. In this section, we slightly modify r -enum and provide three additional data structures, R_{str} , R_{occ} , and R_{eRD} , to recover the original string P and its occurrences in T from the element $\text{output}(P)$. The three data structures R_{str} , R_{occ} and R_{eRD} require $O(rw)$ bits of space and support *extract*, *extract-sa*, and *extended range distinct queries*, respectively. An *extract* query returns string P for a given pair $(\text{interval}(P), |P|)$. An *extract-sa* query returns all the occurrences of P in T (i.e., $\text{SA}[b..e]$) for a given pair $(\text{interval}(P), \text{SA}[b])$, where $\text{interval}(P) = [b, e]$. An *extended range distinct query* $\text{eRD}(L, b, e, \text{SA}[b])$ returns 4-tuple $(c, p_c, q_c, \text{SA}[p_c])$ for each output $(c, p_c, q_c) \in \text{RD}(L, b, e)$ (i.e., $\text{eRD}(L, b, e, \text{SA}[b]) = \{(c, p_c, q_c, \text{SA}[p_c]) \mid (c, p_c, q_c) \in \text{RD}(L, b, e)\}$). We omit the detailed description of the three data structures because each data structure supports the queries using the well-known properties of RLBWT. Formally, the following lemma holds.

► **Lemma 7.** *Let $k = |\text{eRD}(L, b, e, \text{SA}[b])|$. The three data structures R_{str} , R_{occ} , and R_{eRD} of $O(rw)$ bits of space can support *extract*, *extract-sa*, and *extended range distinct queries* in $O(|P| \log \log_w(n/r))$, $O((e - b + 1) \log \log_w(n/r))$, and $O((k + 1) \log \log_w(n/r))$ time, respectively. We can construct the three data structures in $O(n \log \log_w(n/r))$ time and $O(rw)$ bits of working space by processing the RLBWT of T .*

■ **Table 2** Details of dataset. Table details size in megabytes (MB), string length (n), alphabet size (σ), number of runs in BWT (r), and number of maximal repeats (m) for each data.

Data name	Size [MB]	n	σ	r	m
einstein.de.txt	93	92,758,441	118	101,370	79,469
einstein.en.txt	468	467,626,544	140	290,239	352,590
world leaders	47	46,968,181	90	573,487	521,255
influenza	155	154,808,555	16	3,022,822	7,734,058
kernel	258	257,961,616	161	2,791,368	1,786,102
cere	461	461,286,644	6	11,574,641	10,006,383
coreutils	205	205,281,778	237	4,684,460	2,963,022
Escherichia Coli	113	112,689,515	16	15,044,487	12,011,071
para	429	429,265,758	6	15,636,740	13,067,128
100genome	307,705	307,705,110,945	6	36,274,924,494	52,172,752,566

Proof. See Appendix D. ◀

We modify r-enum as follows. The modified r-enum outputs pair $(\text{repr}(P), \text{SA}[b])$, 3-tuple $(\text{interval}(P'), |P'|, \text{SA}[b'])$, and 4-tuple $(\text{interval}(P''), |P''|, c, \text{SA}[b''])$ instead of rich representation $\text{repr}(P)$, pair $(\text{interval}(P'), |P'|)$, and 3-tuple $(\text{interval}(P''), |P''|, c)$, respectively. Here, (i) $P, P', P''c$ are a maximal repeat, minimal unique substring, and minimal absent word, respectively, and (ii) b, b', b'' are the left boundaries of $\text{interval}(P)$, $\text{interval}(P')$, and $\text{interval}(P''c)$, respectively. We replace each range distinct query $\text{RD}(L, b, e)$ used by r-enum with the corresponding extended range distinct query $\text{eRD}(L, b, e, \text{SA}[b])$ to compute the sa-values $\text{SA}[b]$, $\text{SA}[b']$, and $\text{SA}[b'']$. See Appendix E for a detailed description of the modified r-enum. Formally, the following theorem holds.

► **Theorem 8.** *R-enum can also output the sa-values $\text{SA}[b]$, $\text{SA}[b']$, and $\text{SA}[b'']$ for each maximal repeat P , minimal unique substring P' , and minimal absent word $P''c$, respectively, using an additional data structure of $O(rw)$ bits of space. Here, b, b', b'' are the left boundaries of $\text{interval}(P)$, $\text{interval}(P')$, and $\text{interval}(P''c)$, respectively, and the modification does not increase the running time.*

Proof. See Appendix E. ◀

We compute each characteristic substring and all the occurrences of the substring in T by applying extract and extract-sa queries to the corresponding element outputted by the modified r-enum. For example, the outputted pair $(\text{repr}(P), \text{SA}[b])$ for a maximal repeat P contains $\text{interval}(P)$, $|P|$, and $\text{SA}[b]$. Hence, we can obtain P by applying an extract query to pair $(\text{interval}(P), |P|)$. Similarly, we can obtain all the occurrences of P in T by applying an extract-sa query to pair $(\text{interval}(P), \text{SA}[b])$. Note that we do not need to compute the occurrences of minimal absent words in T since the words do not occur in T .

6 Experiments

We demonstrate the effectiveness of our r-enum for enumerating maximal repeats on a benchmark dataset of highly repetitive strings in a comparison with the state-of-the-art enumeration algorithms of the OT [30], BBO [12] and BC [7] methods, which are reviewed in Section 1. In this experiment, all the methods output suffix-tree intervals for maximal repeats.

■ **Table 3** Peak memory consumption of each method in mega bytes (MB).

Data name	Size [MB]	Memory (MB)			
		r-enum	BBO	BC	OT
einstein.de.txt	93	2	100	100	1,642
einstein.en.txt	468	4	488	488	8,278
world leaders	47	5	37	37	832
influenza	155	18	77	73	2,741
kernel	258	21	297	292	4,567
cere	461	86	265	224	8,166
coreutils	205	32	268	237	3,634
Escherichia Coli	113	109	116	55	1,995
para	429	116	262	204	7,599

The OT method enumerates maximal repeats using the BWT and enhanced suffix array of a given string, where the enhanced suffix array consists of suffix and LCP arrays. The OT method runs in $O(n)$ time and with $O(n \log n)$ bits of working space for T .

The BBO method traverses the Weiner-link tree of a given string T using Lemma 1 and a breadth-first search, and it outputs maximal repeats by processing all the nodes in the tree. We used the SDSL library [22] for an implementation of the Huffman-based wavelet tree to support range distinct queries, and we did not implement the technique of Beller et al. for storing a queue for suffix-tree intervals in $n + o(n)$ bits. Hence, our implementation of the BBO method takes the BWT of T and runs in $O(n \log \sigma)$ time and with $|WT_{\text{huff}}| + O(\max\{|\mathcal{K}_0|, |\mathcal{K}_1|, \dots, |\mathcal{K}_n|\})w$ bits of working space, where (i) $|WT_{\text{huff}}| = O(n \log \sigma)$ is the size of the Huffman-based wavelet tree, and (ii) $\mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_n$ are introduced in Section 3.3. The latter term can be bounded by $O(rw)$ bits by applying the analysis described in Section 3.3 to their enumeration algorithm.

The BC method also traverses the Weiner-link tree by Lemma 1 and a depth-first search. The method stores a data structure for range distinct queries and a stack data structure of size $O(\sigma^2 \log^2 n)$ bits for the depth-first search. We also used the SDSL library [22] for the Huffman-based wavelet tree to support range distinct queries. Hence, our implementation of the BC method runs in $O(n \log \sigma)$ time and with $|WT_{\text{huff}}| + O(\sigma^2 \log^2 n)$ bits of working space.

We used a benchmark dataset of nine highly repetitive strings in the Pizza & Chili corpus downloadable from <http://pizzachili.dcc.uchile.cl>. In addition, we demonstrated the scalability of r-enum by enumerating maximal repeats on the concatenation of 100 human genomes (307 gigabytes) built from 1,000 human genomes [33]. Table 2 details our dataset.

We used memory consumption and execution time as evaluation measures for each method. Since each method takes the BWT of a string as an input and outputs suffix-tree intervals for maximal repeats, the execution time consists of two parts: (i) the preprocessing time for constructing data structures built from an input BWT, and (ii) the enumeration time after the data structures are constructed. We performed all the experiments on 48-core Intel Xeon Gold 6126 (2.60 GHz) CPU with 2 TB of memory. The source codes used in the experiments are available at <https://github.com/TNishimoto/renum>.

■ **Table 4** Execution time of each method in seconds (s).

Data name	Execution time (s)			
	r-enum	BBO	BC	OT
einstein.de.txt	172	84	70	13
einstein.en.txt	856	487	387	80
world leaders	97	24	16	10
influenza	267	69	49	30
kernel	559	281	178	57
cere	985	277	186	110
coreutils	445	285	179	42
Escherichia Coli	253	68	40	29
para	961	262	173	102

■ **Table 5** Execution time in hours and peak memory consumption in mega-bytes (MB) of r-enum on 100genome.

Size [MB]	n/r	Execution time (hours)	Memory (MB)
307,705	8.5	25	319,949

6.1 Experimental results on benchmark dataset

In the experiments on the nine highly repetitive benchmark strings, we ran each method and with a single thread. Table 3 shows the peak memory consumption of each method. The BBO and BC methods consumed approximately 1.0-2.0 and 0.9-1.5 bytes per byte of input, respectively. The memory usage of the BC method was no more than that of the BBO method on each of the nine strings. The OT method consumed approximately 18 bytes per character, which was larger than the memory usage of the BBO method. Our r-enum consumed approximately 7-23 bytes per run in BWT. The memory usage of r-enum was the smallest on most of the nine strings except for the file Escherichia Coli. In the best case, the memory usage of r-enum was approximately 122 times less than that of the BC method on einstein.en.txt because the ratio $n/r \approx 1611$ and alphabet size $\sigma = 140$ were large.

Table 4 shows the execution time for each method. The OT method was the fastest among all the methods, and it took approximately 13-110 seconds. The execution times of the BBO and BC methods were competitive, and each execution of them finished within 487 seconds on the nine strings. R-enum was finished in 985 seconds even for the string data (cere) with the longest enumeration time. These results show that r-enum can space-efficiently enumerate maximal substrings in a practical amount of time, although r-enum was approximately 10 times slower than the fastest method (i.e., OT method).

6.2 Experimental results on 100 human genomes

We tested r-enum on 100genome, which is a 307-gigabyte string of 100 human genomes. For this experiment, we implemented computations of Weiner-links from nodes with the same depth in parallel. It is easy to achieve the parallelization, because r-enum uses Lemma 1 to compute Weiner-links and we can apply the Lemma to each node independently. We ran the parallelized r-enum with 48 threads.

Table 5 shows the total execution time and peak memory consumption of r-enum on 100genome. R-enum consumed approximately 25 hours and 319 gigabytes of memory for enumeration. The result demonstrates the scalability and practicality of r-enum for

enumerating maximal repeats on a huge string existing in the real-world. The execution time was 0.6 seconds per 10^6 characters. The memory consumption was almost the same as the size of the input file. It was relatively large compared with the memory consumption of r-enum on the nine benchmark strings because the ratio of $n/r \approx 8.4$ for 100genome was small.

7 Conclusion

We presented r-enum, which can enumerate maximal repeats, minimal unique substrings, and minimal absent words working in $O(rw)$ bits of working space. Experiments using a benchmark dataset of highly repetitive strings showed that r-enum is more space-efficient than the previous methods if $n/r > 27.5$. In addition, we demonstrated the applicability of r-enum to a huge string by performing experiments on a 300-gigabyte string of 100 human genomes. Our method was not so effective for the 300-gigabyte string because the string is weakly compressible. On the other hand, it is obvious that r-enum is more space-efficient than the previous methods for huge strings such that n/r is sufficiently large. Our future work is to reduce the running time and working space of r-enum.

We showed that breadth-first traversal of a Weiner-link tree can be performed in $O(rw)$ bits of working space. The previous breadth-first traversal algorithm by the BBO method requires $|RD| + O(n)$ bits of working space, where $|RD|$ is the size of a data structure supporting range distinct queries on the BWT of T . In addition to enumerations of characteristic substrings, traversal algorithms of a Weiner-link tree can be used for constructing three data structures: (i) LCP array, (ii) suffix tree topology, and (iii) the merged BWT of two strings [32]. Constructing these data structures in $O(rw)$ bits of working space by modifying r-enum would be an interesting future work.

We also showed that the data structure RD' for our traversal algorithm supported two operations: (i) a range distinct query in $O(\log \log_w(n/r))$ time per output element and (ii) computations of LF function in $O(\log \log_w(n/r))$ time. The result can replace the pair $(|RD|, d)$ presented in Table 1 with pair $(O(rw), O(\log \log_w(n/r)))$, which improves the BBO and BC methods so that they work in $O(rw + n)$ and $O(rw + \sigma^2 \log^2 n)$ bits of working space, respectively. We think that the working space of the BC method with RD' is practically smaller than that of r-enum, because $\sigma^2 \log^2 n$ is practically smaller than rw in many cases. This insight indicates that we can enumerate characteristic substrings with a lower memory consumption of $O(rw)$ bits by combining r-enum with the BC method even for a large alphabet size. Thus, the following method could improve the space efficiency for enumerations: executing the BC method with RD' for a small alphabet (i.e., $\sigma < \sqrt{r/\log n}$) and executing r-enum for a large alphabet.

References

- 1 Paniz Abedin, M. Oguzhan Külekci, and Shama V. Thankachan. A survey on shortest unique substring queries. *Algorithms*, 13:224, 2020.
- 2 Mohamed Ibrahim Abouelhoda, Stefan Kurtz, and Enno Ohlebusch. Replacing suffix trees with enhanced suffix arrays. *Journal of Discrete Algorithms*, 2:53–86, 2004.
- 3 Alberto Apostolico. The myriad virtues of subword trees. In *Combinatorial Algorithms on Words*, pages 85–96, 1985.
- 4 Hideo Bannai, Travis Gagie, and Tomohiro I. Refining the r -index. *Theoretical Computer Science*, 812:96–108, 2020.
- 5 Carl Barton, Alice Héliou, Laurent Mouchard, and Solon P. Pissis. Linear-time computation of minimal absent words using suffix array. *BMC Bioinformatics*, 15:388, 2014.

- 6 Verónica Becher, Alejandro Deymonnaz, and Pablo Ariel Heiber. Efficient computation of all perfect repeats in genomic sequences of up to half a gigabyte, with a case study on the human genome. *Bioinformatics*, 25:1746–1753, 2009.
- 7 Djamel Belazzougui and Fabio Cunial. Space-efficient detection of unusual words. In *Proceedings of SPIRE*, pages 222–233, 2015.
- 8 Djamel Belazzougui, Fabio Cunial, Travis Gagie, Nicola Prezza, and Mathieu Raffinot. Composite repetition-aware data structures. In *Proceedings of CPM*, pages 26–39, 2015.
- 9 Djamel Belazzougui, Fabio Cunial, Juha Kärkkäinen, and Veli Mäkinen. Linear-time string indexing and analysis in small space. *ACM Transactions on Algorithms*, 16:17:1–17:54, 2020.
- 10 Djamel Belazzougui and Gonzalo Navarro. Optimal lower and upper bounds for representing sequences. *ACM Transactions on Algorithms*, 11:31:1–31:21, 2015.
- 11 Djamel Belazzougui, Gonzalo Navarro, and Daniel Valenzuela. Improved compressed indexes for full-text document retrieval. *Journal of Discrete Algorithms*, 18:3–13, 2013.
- 12 Timo Beller, Katharina Berger, and Enno Ohlebusch. Space-efficient computation of maximal and supermaximal repeats in genome sequences. In *Proceedings of SPIRE*, pages 99–110, 2012.
- 13 Michael Burrows and David J Wheeler. A block-sorting lossless data compression algorithm. *Technical report*, 1994.
- 14 Panagiotis Charalampopoulos, Maxime Crochemore, Gabriele Fici, Robert Mercas, and Solon P. Pissis. Alignment-free sequence comparison using absent words. *Information and Computation*, 262:57–68, 2018.
- 15 Francisco Claude, Gonzalo Navarro, and Alberto Ordóñez Pereira. The wavelet matrix: An efficient wavelet tree for large alphabets. *Information Systems*, 47:15–32, 2015.
- 16 Maxime Crochemore, Gabriele Fici, Robert Mercas, and Solon P. Pissis. Linear-time sequence comparison using minimal absent words & applications. In *Proceedings of LATIN*, pages 334–346, 2016.
- 17 Maxime Crochemore, Filippo Mignosi, and Antonio Restivo. Automata and forbidden words. *Information Processing Letters*, 67:111–117, 1998.
- 18 Maxime Crochemore, Filippo Mignosi, Antonio Restivo, and Sergio Salemi. Data compression using antidictionaries. *Proceedings of the IEEE*, 88:1756–1768, 2000.
- 19 Maxime Crochemore and Gonzalo Navarro. Improved antidictionary based compression. In *Proceedings of SCCC*, pages 7–13, 2002.
- 20 Isamu Furuya, Takuya Takagi, Yuto Nakashima, Shunsuke Inenaga, Hideo Bannai, and Takuya Kida. MR-RePair: Grammar compression based on maximal repeats. In *Proceedings of DCC*, pages 508–517, 2019.
- 21 Travis Gagie, Gonzalo Navarro, and Nicola Prezza. Fully functional suffix trees and optimal text searching in bwt-runs bounded space. *Journal of the ACM*, 67:2:1–2:54, 2020.
- 22 Simon Gog, Timo Beller, Alistair Moffat, and Matthias Petri. From theory to practice: Plug and play with succinct data structures. In *Proceedings of SEA*, pages 326–337, 2014.
- 23 Dan Gusfield. *Algorithms on Strings, Trees, and Sequences - Computer Science and Computational Biology*. Cambridge University Press, 1997.
- 24 Bernhard Haubold, Nora Pierstorff, Friedrich Möller, and Thomas Wiehe. Genome comparison without alignment using shortest unique substrings. *BMC Bioinformatics*, 6:123, 2005.
- 25 Dominik Kempa. Optimal construction of compressed indexes for highly repetitive texts. In *Proceedings of SODA*, pages 1344–1357, 2019.
- 26 Dominik Kempa and Tomasz Kociumaka. Resolution of the burrows-wheeler transform conjecture. In *Proceedings of FOCS*, pages 1002–1013, 2020.
- 27 Udi Manber and Eugene W. Myers. Suffix arrays: A new method for on-line string searches. *SIAM Journal on Computing*, 22:935–948, 1993.
- 28 Tomonari Masada, Atsuhiko Takasu, Yuichiro Shibata, and Kiyoshi Oguri. Clustering documents with maximal substrings. In *Proceedings of ICEIS*, pages 19–34, 2011.

- 29 Tatsuya Ohno, Kensuke Sakai, Yoshimasa Takabatake, Tomohiro I, and Hiroshi Sakamoto. A faster implementation of online RLBWT and its application to LZ77 parsing. *Journal of Discrete Algorithms*, 52-53:18–28, 2018.
- 30 Daisuke Okanohara and Jun'ichi Tsujii. Text categorization with all substring features. In *Proceedings of SDM*, pages 838–846, 2009.
- 31 Alberto Policriti and Nicola Prezza. LZ77 computation based on the run-length encoded BWT. *Algorithmica*, 80:1986–2011, 2018.
- 32 Nicola Prezza and Giovanna Rosone. Space-efficient construction of compressed suffix trees. *Theoretical Computer Science*, 852:138–156, 2021.
- 33 The 1000 Genomes Project Consortium. An integrated map of genetic variation from 1,092 human genomes. *Nature*, 491:56–65, 2012.

A Algorithm for constructing data structures in Section 3.1

Recall that our traversal algorithm uses six data structures: (i) the RLBWT of T , (ii) $\text{repr}(\varepsilon)$, (iii) D_{LF} , (iv) $R_{\text{rank}}(\mathcal{S}_{\text{start}})$, (v) $R_{\text{RD}}(L')$, and (vi) an empty array X of size σ . Let δ be the permutation of $[1, r]$ satisfying either of two conditions for two distinct integers $i, j \in \{1, 2, \dots, r\}$: (i) $L[\ell(\delta[i])] \prec L[\ell(\delta[j])]$ or (ii) $L[\ell(\delta[i])] = L[\ell(\delta[j])]$ and $i < j$. Then $D_{\text{LF}}[1] = 1$ and $D_{\text{LF}}[i] = D_{\text{LF}}[i-1] + |L[\ell(\delta[i-1]).. \ell(\delta[i-1]+1)] - 1|$ hold by LF formula. We construct the permutation δ in $O(n)$ time and $O(rw)$ bits of working space using LSD radix sort. After that we construct the array D_{LF} in $O(r)$ time using δ .

Next, we construct $\mathcal{S}_{\text{start}}$ and L' in $O(r)$ time by processing the RLBWT of T . After that, we construct $R_{\text{rank}}(\mathcal{S}_{\text{start}})$ in $O(|\mathcal{S}_{\text{start}}| \log \log_w(n/r))$ time and $O(|\mathcal{S}_{\text{start}}|w)$ bits of working space by processing $\mathcal{S}_{\text{start}}$ [21, Appendix A.1]. Similarly, we construct $R_{\text{RD}}(L')$ in $O(r)$ time and $O(r \log \sigma)$ bits of working space [9, Lemma 3.17].

Finally, $\text{repr}(\varepsilon)$ consists of 3-tuple $([1, n], \text{RD}(L, 1, n), 0)$. We compute $\text{RD}(L, 1, n)$ by Lemma 2. Hence the construction time is $O(n \log \log_w(n/r))$ in total, and the working space is $O(rw)$ bits.

B Proof of Theorem 5 (ii)

For simplicity, we focus on minimal unique substrings with a length of at least 2. Every minimal unique substring with a length of at least 2 is a substring cPc' of T , where c, c' are characters, and P is a string with a length of at least 0. R-enum uses an array Right_P of size σ for detecting a minimal unique substring cPc' . $\text{Right}_P[c]$ stores $|\text{Occ}(T, Pc)|$ for each $c \in \Sigma$. A substring cPc' of T is a minimal unique substring if and only if cPc' satisfies three conditions: (i) $|\text{Occ}(T, cPc')| = 1$, (ii) $|\text{Occ}(T, cP)| \geq 2$, and (iii) $\text{Right}_P[c] \geq 2$ hold. We can verify the three conditions using $\text{repr}(P)$ and Lemma 1, and hence, the following lemma holds.

► **Lemma 9.** *Let $\mathcal{M}(P)$ be the set of minimal unique substrings such that the form of each minimal unique substring is cPc' , where c, c' are characters, and P is a given string. We can compute the output by r-enum for the set $\mathcal{M}(P)$ (i.e., $\{(\text{interval}(cPc'), |cPc'|) \mid cPc' \in \mathcal{M}(P)\}$) using (i) $\text{repr}(R)$, (ii) the data structures presented in Section 3.1, and (iii) an empty array X' of size σ . The running time and working space are $O(h \log \log_w(n/r))$ and $O((\sigma + h')w)$ bits, respectively. Here, $h = \sum_{Pc' \in \text{children}(P)} |\text{WLink}(Pc')|$, and $h' = \sum_{cP \in \text{WLink}(P) \cap \mathcal{L}_{|P|+1}} |\text{children}(cP)|$.*

Proof. $\text{Right}_P[c]$ is stored in $\text{repr}(P)$ for $c \in \Sigma$, and the pair $(\text{interval}(cPc'), |cPc'|)$ is stored in $\text{repr}(cP)$. We compute the output by four steps. (i) We convert X' into Right_P by processing $\text{repr}(P)$. (ii) We compute the rich representations for all the strings in $\text{WLink}(P)$ by Lemma 1.

(iii) We process the rich representation for each string $cP \in \text{WLink}(P)$ and output pair $(\text{interval}(cPc'), |cPc'|)$ for each child cPc' if $|\text{interval}(cPc')| = 1$ and $\text{Right}_P[c'] \geq 2$. (iv) We recover X' from Right_P . Hence, Lemma 9 holds. \blacktriangleleft

Node u_P is always explicit (i.e., u_P is a node of the Weiner-link tree) if a substring cPc' is a minimal unique substring of T because $|\text{Occ}(T, P)| > |\text{Occ}(T, Pc')|$ holds from the definition of the minimal unique substring. Hence, we can compute $(\text{interval}(cPc'), |cPc'|)$ for each minimal unique substring cPc' in T by applying Lemma 9 to all the nodes in the Weiner-link tree.

R-enum prepares an empty array X' of size σ and computes the output for set $\mathcal{M}(P)$ by applying Lemma 9 to each node u_P . The running time and working space are $O((\sum_{t=0}^n H_t) \log \log_w(n/r)) = O(n \log \log_w(n/r))$ and $O((r + \sigma + \max\{|\mathcal{K}_0|, |\mathcal{K}_1|, \dots, |\mathcal{K}_n|\})w) = O(rw)$ bits, respectively. Here, H_t and \mathcal{K}_t are the terms introduced in Section 3.3.

C Proof of Theorem 5 (iii)

For simplicity, we focus on minimal absent words with a length of at least 2. Minimal absent words have similar properties to the properties of the minimal unique substrings explained in the proof of Theorem 5(ii), i.e., the characteristic substrings have two properties: (i) a string cPc' is a minimal absent word for T if and only if cPc' satisfies three conditions: (1) $|\text{Occ}(T, cPc')| = 0$, (2) $|\text{Occ}(T, cP)| \geq 1$, and (3) $\text{Right}_P[c'] \geq 1$ hold, and (ii) node u_P is always explicit if a substring cPc' is a minimal absent word for T because $|\text{Occ}(T, P)| > |\text{Occ}(T, Pc')|$ holds from the definition of the minimal absent word. We obtain the following lemma by modifying the proof of Lemma 9.

► **Lemma 10.** *Let $\mathcal{W}(P)$ be the set of minimal absent words such that the form of each minimal unique substring is cPc' , where c, c' are characters, and P is a given string. We can compute the output by r-enum for the set $\mathcal{W}(P)$ (i.e., $\{(\text{interval}(cP), |cP|, c') \mid cPc' \in \mathcal{W}(P)\}$) using (i) $\text{repr}(R)$, (ii) the data structures presented in Section 3.1, and (iii) an empty array X' of size σ . The running time and working space are $O(h \log \log_w(n/r) + |\mathcal{W}(P)|)$ and $O((\sigma + h')w)$ bits, respectively.*

We can compute $(\text{interval}(cP), |cP|, c')$ for each minimal absent word cPc' for T by applying Lemma 10 to all the nodes in the Weiner-link tree. R-enum prepares an empty array X' of size σ and computes the output for set $\mathcal{W}(P)$ by applying Lemma 10 to each node u_P . The running time and working space are $O(n \log \log_w(n/r) + occ)$ and $O(rw)$ bits, respectively.

D Proof of Lemma 7

Proof for data structure R_{str} . Data structure R_{str} consists of (i) the RLBWT of T , (ii) π , (iii) D_{LF} , (iv) $R_{\text{rank}}(\mathcal{S}'_{\text{start}})$, and (v) $R_{\text{rank}}(\mathcal{S}_{\text{start}})$. Here, (i) π is the permutation on $\{1, 2, \dots, r\}$ satisfying $\text{LF}(\ell(\pi[1])) < \text{LF}(\ell(\pi[2])) < \dots < \text{LF}(\ell(\pi[r]))$, (ii) D_{LF} is the array introduced in Section 3, and (iii) $\mathcal{S}'_{\text{start}} = \{\text{LF}(\ell(\pi[1])), \text{LF}(\ell(\pi[2])), \dots, \text{LF}(\ell(\pi[r]))\}$. $\pi[i] = \delta[i]$ holds for any integer $i \in \{1, 2, \dots, r\}$. $\mathcal{S}'_{\text{start}}$ and D_{LF} can be constructed in $O(r)$ time after the permutation π is constructed. We already showed that δ could be constructed in $O(n)$ time by processing the RLBWT of T , which was explained in Appendix A. Hence, we can construct R_{str} in $O(n \log \log_w(n/r))$ time and $O(rw)$ bits of working space.

We introduce the inverse function LF^{-1} of LF function (i.e., $\text{LF}^{-1}(\text{LF}(i)) = i$ holds for $i \in \{1, 2, \dots, n\}$) to solve the extract query. Recall that $\text{LF}(i) = D_{\text{LF}}[x] + (i - \ell(x))$ holds, which is shown in Section 3.1, where $x = \text{rank}(\mathcal{S}_{\text{start}}, i)$. Similarly, $\text{LF}^{-1}(i) = \ell(\pi[y]) + (i - D_{\text{LF}}[\pi[y]])$ holds by the LF formula for any integer $i \in \{1, 2, \dots, n\}$, where $y = \text{rank}(\mathcal{S}'_{\text{start}}, i)$. Hence, we can compute $\text{LF}^{-1}(i)$ in $O(\log \log_w(n/r))$ time using the data structure R_{str} .

Let LF_x^{-1} be the function that returns the position obtained by recursively applying the inverse LF function to i x times (i.e., $\text{LF}_0^{-1}(i) = i$, and $\text{LF}_x^{-1}(i) = \text{LF}^{-1}(\text{LF}_{x-1}^{-1}(i))$). Then, $T[\text{SA}[i].. \text{SA}[i] + d - 1] = L[\text{LF}_1^{-1}(i)], L[\text{LF}_2^{-1}(i)], \dots, L[\text{LF}_d^{-1}(i)]$ holds for any integer $d \geq 1$ because $\text{SA}[\text{LF}^{-1}(i)] = \text{SA}[i] + 1$ holds unless $\text{SA}[i] = n$. R_{str} can support random access to the BWT L in $O(\log \log_w(n/r))$ time using a rank query on set $\mathcal{S}_{\text{start}}$. Hence, we can compute $T[\text{SA}[i].. \text{SA}[i] + d - 1]$ in $O(d \log \log_w(n/r))$ time using R_{str} for two given integers i and d .

We explain an algorithm solving an extract query for a given rich representation $\text{repr}(P)$. Let $\text{interval}(P) = [b, e]$. Then, $\text{SA}[b]$ stores the index of a suffix of T having P as a prefix, i.e., $T[\text{SA}[b].. \text{SA}[b] + |P| - 1] = P$ holds. We recover the prefix P from $\text{SA}[b]$ using R_{str} . Hence R_{str} supports the extract query in $O(|P| \log \log_w(n/r))$ time. ◀

Proof for data structure R_{occ} . Next, we leverage a function ϕ to solve the extract-sa query. The function $\phi(\text{SA}[i])$ returns $\text{SA}[i + 1]$ for $i \in \{1, 2, \dots, n - 1\}$. R_ϕ is a data structure of $O(rw)$ bits proposed by Gagie et al. [21], and we can compute ϕ function in $O(\log \log_w(n/r))$ time by R_ϕ . The data structure can be constructed in $O(n \log \log_w(n/r))$ time and $O(rw)$ bits of working space by processing the RLBWT of T [21]. The second data structure R_{occ} consists of R_ϕ , and we solve the extract-sa query by recursively applying the function ϕ to $\text{SA}[b]$ $(e - b)$ times. Hence R_{occ} can support the extract-sa query in $O((e - b + 1) \log \log_w(n/r))$ time. ◀

Proof for data structure R_{eRD} . Let D_{SA} be an array of size r such that $D_{\text{SA}}[i]$ stores the sa-value at the starting position of the i -th run in BWT L for $i \in \{1, 2, \dots, r\}$, i.e., $D_{\text{SA}}[i] = \text{SA}[\ell(i)]$. Let $(c, p_c, q_c, \text{SA}[p_c])$ be a 4-tuple outputted by query $\text{eRD}(L, b, e, \text{SA}[b])$ and x be the index of the run containing character $L[p_c]$ (i.e., $x = \text{rank}(\mathcal{S}_{\text{start}}, p_c)$). Then $\text{SA}[p_c] = D_{\text{SA}}[x]$ if $p_c = \ell(x)$; otherwise $\text{SA}[p_c] = \text{SA}[b]$ holds because p_c is equal to $\ell(x)$ or b under Lemma 2. The relationship among $\text{SA}[p_c]$, $D_{\text{SA}}[x]$, and $\text{SA}[b]$ is called the *toehold lemma* (e.g., [31, 21]). The toehold lemma indicates that we can compute $\text{SA}[p_c]$ in $O(\log \log_w(n/r))$ time for each output $(c, p_c, q_c) \in \text{RD}(L, b, e)$ using (i) the array D_{SA} , (ii) data structure $R_{\text{rank}}(\mathcal{S}_{\text{start}})$, and (iii) the RLBWT of T if we know the first sa-value $\text{SA}[b]$ in $[b, e]$.

Next, we explain R_{eRD} . R_{eRD} consists of $R_{\text{RD}}(L')$, D_{SA} , $R_{\text{rank}}(\mathcal{S}_{\text{start}})$, and the RLBWT of T . We already showed that we could construct $R_{\text{RD}}(L')$ and $R_{\text{rank}}(\mathcal{S}_{\text{start}})$ in $O(n \log \log_w(n/r))$ time by processing the RLBWT of T . We construct the array D_{SA} by computing all the sa-values in $\text{SA}[1..n]$ in left-to-right order using data structure R_ϕ . R_ϕ can be constructed in $O(n \log \log_w(n/r))$ time by processing the RLBWT of T . Hence, the construction time for R_{eRD} is $O(n \log \log_w(n/r))$ time in total, and the working space is $O(rw)$ bits.

We solve extended range distinct query $\text{eRD}(L, b, e, \text{SA}[b])$ using the toehold lemma after solving the corresponding range distinct query $\text{RD}(L, b, e)$ using $R_{\text{RD}}(L')$, $R_{\text{rank}}(\mathcal{S}_{\text{start}})$, and the RLBWT of T . The running time is $O((k + 1) \log \log_w(n/r))$, where $k = |\text{eRD}(L, b, e, \text{SA}[b])|$. ◀

E Proof of Theorem 8

We extend Lemma 1. Let $\text{eRepr}(P)$ for P be a 4-tuple $(\text{interval}(P), \{(c_1, b_1, e_1, \text{SA}[b_1]), (c_2, b_2, e_2, \text{SA}[b_2]), \dots, (c_k, b_k, e_k, \text{SA}[b_k])\}, |P|, \text{SA}[b])$. Here, (i) b is the left boundary of $\text{interval}(P)$, (ii) Pc_1, Pc_2, \dots, Pc_k are strings represented by the children of node u_P , and (iii) $[b_i, e_i] = \text{interval}(Pc_i)$ for $i \in [1, k]$. We call $\text{eRepr}(P)$ an *extended rich representation*.

Let $\text{eRepr}(cP) = (\text{interval}(cP), \{(c'_1, b'_1, e'_1, \text{SA}[b'_1]), (c'_2, b'_2, e'_2, \text{SA}[b'_2]), \dots, (c'_{k'}, b'_{k'}, e'_{k'}, \text{SA}[b'_{k'}])\}, |cP|, \text{SA}[b'])$ for a character c . Let x_i be an integer such that $\text{LF}(x_i) = b'_i$ for a 4-tuple $(c'_i, b'_i, e'_i, \text{SA}[b'_i])$ in $\text{eRepr}(P)$, and let $y(i)$ be an integer such that $x_i \in [b_{y(i)}, e_{y(i)}]$

holds. Then, there exists a tuple $(\hat{c}, p_{\hat{c}}, q_{\hat{c}}, \text{SA}[p_{\hat{c}}]) \in \text{eRD}(L, b_{y(i)}, e_{y(i)}, \text{SA}[b_{y(i)}])$ such that $p_{\hat{c}} = x_i$ holds. $\text{SA}[b'_i] = \text{SA}[p_{\hat{c}}] - 1$ holds by LF function. Next, let j be an integer such that b'_j is the smallest in set $\{b'_1, b'_2, \dots, b'_{k'}\}$. Then $\text{SA}[b'] = \text{SA}[b'_j]$ holds because b'_j is equal to the left boundary of $\text{interval}(cP)$. Hence, Lemma 1 can output extended rich representations instead of rich representations by replacing the range distinct queries used by the algorithm of Lemma 1 with the corresponding extended range distinct queries. Formally, the following lemma holds.

► **Lemma 11.** *We can compute set $\{\text{eRepr}(cP) \mid cP \in \text{WLink}(P)\}$ for a given rich representation $\text{eRepr}(P)$ in $O(h \log \log_w(n/r))$ time using R_{eRD} and the six data structures introduced in Section 3.1, where $h = \sum_{Pc' \in \text{children}(P)} |\text{WLink}(Pc')|$.*

Next, we modify our traversal algorithm for the Weiner-link tree of T . The modified traversal algorithm uses Lemma 11 instead of Lemma 1. Hence, we obtain the following lemma.

► **Lemma 12.** *We can output n sets $\{\text{eRepr}(P) \mid P \in \mathcal{L}_0\}$, $\{\text{eRepr}(P) \mid P \in \mathcal{L}_1\}$, \dots , $\{\text{eRepr}(P) \mid P \in \mathcal{L}_{n-1}\}$ in left-to-right order in $O(n \log \log_w(n/r))$ time and $O(rw)$ bits of working space by processing the RLBWT of T .*

Finally, we prove Theorem 8 using the modified traversal algorithm, i.e., Lemma 12.

Proof for maximal repeats. The node u_P representing a maximal repeat P is explicit, and hence, $\text{eRepr}(P)$ is outputted by the modified traversal algorithm. R-enum uses the modified traversal algorithm instead of our original traversal algorithm. Hence, r-enum can output the extended rich representations for all the maximal repeats in T without increasing the running time. ◀

Proof for minimal unique substrings. Let cPc' be a minimal unique substring of T such that its occurrence position is $\text{SA}[b']$. Recall that r-enum computes $\text{repr}(cP)$ by applying Lemma 1 to $\text{repr}(P)$ and outputs $(\text{interval}(cPc'), |cPc'|)$ by processing $\text{repr}(cP)$. The extended rich representation $\text{eRepr}(cP)$, which corresponds to $\text{repr}(cP)$, contains $\text{interval}(cPc')$, $|cPc'|$, and $\text{SA}[b']$. The modified r-enum (i) traverses the Weiner-link tree by the modified traversal algorithm, (ii) computes $\text{eRepr}(cP)$ by applying Lemma 11 to $\text{eRepr}(P)$, and (iii) outputs $(\text{interval}(cPc'), |cPc'|, \text{SA}[b'])$ by processing $\text{eRepr}(cP)$. The running time is $O(n \log \log_w(n/r))$ time in total. ◀

Proof for minimal absent words. Let cPc' be a minimal absent word for T , and let b'' be the left boundary of $\text{interval}(cP)$. $\text{eRepr}(cP)$ contains the sa-value $\text{SA}[b'']$, and hence, we can compute $(\text{interval}(cP), |cP|, c', \text{SA}[b'])$ for the minimal absent word cPc' by modifying the algorithm used by the modified r-enum for minimal unique substrings.

The modified r-enum (i) traverses the Weiner-link tree by the modified traversal algorithm, (ii) computes $\text{eRepr}(cP)$ by applying Lemma 11 to $\text{eRepr}(P)$, and (iii) outputs $(\text{interval}(cP), |cP|, c', \text{SA}[b'])$ by processing $\text{eRepr}(cP)$. The running time is $O(n \log \log_w(n/r) + \text{occ})$ time in total. ◀

F Omitted table

■ **Table 6** Execution time of each method. Execution time is separately presented as enumeration and preprocessing times in seconds (s).

Data name	Preprocessing time [s]				Enumeration time [s]			
	r-enum	BBO	BC	OT	r-enum	BBO	BC	OT
einstein.de.txt	1	1	1	7	171	83	69	6
einstein.en.txt	3	3	3	50	853	484	384	30
world leaders	1	1	1	7	96	23	15	3
influenza	2	1	1	17	265	68	48	13
kernel	2	2	2	41	557	279	176	16
cere	5	3	3	78	980	274	183	32
coreutils	2	2	2	29	443	283	177	13
Escherichia Coli	5	1	1	18	248	67	39	11
para	6	3	3	71	955	259	170	31