# A framework for space-efficient string kernels 

Djamal Belazzougui ${ }^{1,2}$ and Fabio Cunial ${ }^{1,2}$<br>${ }^{1}$ Department of Computer Science, University of Helsinki, Finland.*<br>${ }^{2}$ Helsinki Institute for Information Technology, Finland.

September 6, 2021


#### Abstract

String kernels are typically used to compare genome-scale sequences whose length makes alignment impractical, yet their computation is based on data structures that are either space-inefficient, or incur large slowdowns. We show that a number of exact string kernels, like the $k$-mer kernel, the substrings kernels, a number of length-weighted kernels, the minimal absent words kernel, and kernels with Markovian corrections, can all be computed in $O(n d)$ time and in $o(n)$ bits of space in addition to the input, using just a rangeDistinct data structure on the Burrows-Wheeler transform of the input strings, which takes $O(d)$ time per element in its output. The same bounds hold for a number of measures of compositional complexity based on multiple value of $k$, like the $k$-mer profile and the $k$-th order empirical entropy, and for calibrating the value of $k$ using the data.


## 1 Introduction

Given two strings $T^{1}$ and $T^{2}$, a kernel is a function that simultaneously converts $T^{1}$ and $T^{2}$ into vectors $\mathbf{T}^{1}$ and $\mathbf{T}^{2}$ in $\mathbb{R}^{n}$ for some $n>0$, and computes a similarity or a distance measure between $\mathbf{T}^{1}$ and $\mathbf{T}^{2}$, without building and storing $\mathbf{T}^{i}$ explicitly [14. Kernels are often the method of choice for comparing extremely long strings, like genomes, read sets, and metagenomic samples, whose size makes alignment infeasible, yet their computation is typically based on space-inefficient data structures, like (truncated) suffix trees, or on space-efficient data structures with large slowdowns, like compressed suffix trees (see e.g. [1, 9] and references therein). The (possibly infinite) dimensions of $\mathbf{T}^{i}$ are, for example, all strings of a specific family on the alphabet of $T^{1}$ and $T^{2}$, and the value assigned to vector $\mathbf{T}^{i}$ along dimension $W$ corresponds to the number of occurrences of string $W$ in $T^{i}$, often rescaled and corrected in domain-specific ways. $\mathbf{T}^{i}$ is often called composition vector, and a large number of its components can be zero in practice. In this paper we focus on space- and time-efficient algorithms for computing the cosine of the angle between two composition vectors $\mathbf{T}^{1}$ and $\mathbf{T}^{2}$, i.e. on computing the kernel $\kappa\left(\mathbf{T}^{1}, \mathbf{T}^{2}\right)=N / \sqrt{D^{1} D^{2}} \in[-1 . .1]$, where $N=\sum_{W} \mathbf{T}^{1}[W] \mathbf{T}^{2}[W]$ and $D^{i}=\sum_{W} \mathbf{T}^{i}[W]^{2}$. This measure of similarity can be converted into a distance $d\left(\mathbf{T}^{1}, \mathbf{T}^{2}\right)=\left(1-\kappa\left(\mathbf{T}^{1}, \mathbf{T}^{2}\right)\right) / 2 \in[0 . .1]$, and the algorithms we describe can be applied to compute norms of vector $\mathbf{T}^{1}-\mathbf{T}^{2}$, like the $p$-norm and the infinity norm. When $\mathbf{T}^{1}$ and $\mathbf{T}^{2}$ are bitvectors, we are more interested in interpreting them as sets and in computing the Jaccard distance $J\left(\mathbf{T}^{1}, \mathbf{T}^{2}\right)=\left\|\mathbf{T}^{1} \wedge \mathbf{T}^{2}\right\| /\left\|\mathbf{T}^{1} \vee \mathbf{T}^{2}\right\|=$ $\left\|\mathbf{T}^{1} \wedge \mathbf{T}^{2}\right\| /\left(\left\|\mathbf{T}^{1}\right\|+\left\|\mathbf{T}^{2}\right\|+\left\|\mathbf{T}^{1} \wedge \mathbf{T}^{2}\right\|\right)$, where $\wedge$ and $\vee$ are the bitwise AND and OR operators, and where $\|\cdot\|$ measures the number of ones in a bitvector.

Given a data structure that supports rangeDistinct queries on the Burrows-Wheeler transform of each string in input, we show that a number of popular string kernels, like the $k$-mer kernel, the substrings kernels, a number of length-weighted kernels, the minimal absent words kernel, and kernels with Markovian

[^0]corrections, can all be computed in $O(n d)$ time and in $o(n)$ bits of space in addition to the input, all in a single pass over the BWTs of the input strings, where $d$ is the time taken by the rangeDistinct query per element in its output. The same bounds hold for computing a number of measures of compositional complexity for multiple values of $k$ at the same time, like the $k$-mer profile and the $k$-th order empirical entropy, and for choosing the value of $k$ used in $k$-mer kernels from the data. All these algorithms become $O(n)$ using the rangeDistinct data structure described in [4], and concatenating this setup to the BWT construction algorithm described in [3], we can compute all such kernels and complexity measures from the input strings in randomized $O(n)$ time and in $O(n \log \sigma)$ bits of space in addition to the input. Finally, we show that measures of expectation based on Markov models are related to the left and right extensions of maximal repeats.

## 2 Preliminaries

### 2.1 Strings

Let $\Sigma=[1 . . \sigma]$ be an integer alphabet, let $\#=0, \#_{1}=-1$ and $\#_{2}=-2$ be distinct separators not in $\Sigma$, and let $T=[1 . . \sigma]^{n-1} \#$ be a string. We assume $\sigma \in o(\sqrt{n} / \log n)$ throughout the paper. A $k$-mer is any string $W \in[1 . . \sigma]$ of length $k>0$. We denote by $f_{T}(W)$ the number of (possibly overlapping) occurrences of a string $W$ in the circular version of $T$, and we use the shorthand $p_{T}(W)=f_{T}(W) /(n-|W|)$ to denote an approximation of the empirical probability of observing $W$ in $T$, assuming that all positions of $T$ except the last $|W|$ ones are equally probable starting positions for $W$. A repeat $W$ is a string that satisfies $f_{T}(W)>1$. We denote by $\Sigma_{T}^{\ell}(W)$ the set of characters $\left\{a \in[0 . . \sigma]: f_{T}(a W)>0\right\}$ and by $\Sigma_{T}^{r}(W)$ the set of characters $\left\{b \in[0 . . \sigma]: f_{T}(W b)>0\right\}$. A repeat $W$ is right-maximal (respectively, left-maximal) iff $\left|\Sigma_{T}^{r}(W)\right|>1$ (respectively, iff $\left|\Sigma_{T}^{\ell}(W)\right|>1$ ). It is well known that $T$ can have at most $n-1$ right-maximal substrings and at most $n-1$ left-maximal substrings. A maximal repeat of $T$ is a repeat that is both left- and right-maximal.

For reasons of space we assume the reader to be familiar with the notion of suffix tree $\mathrm{ST}_{T}$ of a string $T$, and with the notion of generalized suffix tree of two strings, which we do not define here. We denote by $\ell(v)$ the string label of a node $v$ in a suffix tree. It is well known that a substring $W$ of $T$ is right-maximal iff $W=\ell(v)$ for some internal node $v$ of $\mathrm{ST}_{T}$. We assume the reader to be familiar with the notion of suffix link connecting a node $v$ with $\ell(v)=a W$ for some $a \in[0 . . \sigma]$ to a node $w$ with $\ell(w)=W$ : we say that $w=\operatorname{suffixLink}(v)$ in this case. Here we just recall that suffix links and internal nodes of $\mathrm{ST}_{T}$ form a tree, called the suffix-link tree of $T$ and denoted by $\mathrm{SLT}_{T}$, and that inverting the direction of all suffix links yields the so-called explicit Weiner links. Given an internal node $v$ and a symbol $a \in[0 . . \sigma]$, it might happen that string $a \ell(v)$ does occur in $T$, but that it is not right-maximal, i.e. it is not the label of any internal node of $\mathrm{ST}_{T}$ : all such left extensions of internal nodes that end in the middle of an edge are called implicit Weiner links. An internal node $v$ of $\mathrm{ST}_{T}$ can have more than one outgoing Weiner link, and all such Weiner links have distinct labels: in this case, $\ell(v)$ is a maximal repeat. It is known that the number of suffix links (or, equivalently, of explicit Weiner links) is upper-bounded by $2 n-2$, and that the number of implicit Weiner links can be upper-bounded by $2 n-2$ as well.

### 2.2 Enumerating right-maximal substrings and maximal repeats

For reasons of space we assume the reader to be familiar with the notion and uses of the Burrows-Wheeler transform of $T$, including the $C$ array, the rank function, and backward searching. In this paper we use $\mathrm{BWT}_{T}$ to denote the BWT of $T$, we use range $(W)=[\operatorname{sp}(W) . . \mathrm{ep}(W)]$ to denote the lexicographic interval of a string $W$ in a BWT that is implicit from the context, and we use $\Sigma_{i, j}$ to denote the set of distinct characters that occur inside interval $[i . . j]$ of a string that is implicit from the context. We also denote by rangeDistinct $(i, j)$ the function that returns the set of tuples $\left\{\left(c, \operatorname{rank}\left(c, p_{c}\right), \operatorname{rank}\left(c, q_{c}\right)\right): c \in \Sigma_{i, j}\right\}$, where $p_{c}$ and $q_{c}$ are the first and the last occurrence of $c$ inside interval $[i . . j]$, respectively. Here we focus on a specific application of $\mathrm{BWT}_{T}$ : enumerating all the right-maximal substrings of $T$, or equivalently all the internal nodes of $\mathrm{ST}_{T}$. In particular, we use the algorithm described in [3] (Section 4.1), which we sketch
here for completeness.
Given a substring $W$ of $T$, let $b_{1}<b_{2}<\cdots<b_{k}$ be the sorted sequence of all the distinct characters in $\Sigma^{r}(W)$, and let $a_{1}, a_{2}, \ldots, a_{h}$ be the list of all the characters in $\Sigma^{\ell}(W)$, not necessarily sorted. Assume that we represent a substring $W$ of $T$ as a pair repr $(W)=($ chars $[1 . . k]$, first $[1 . . k+1]$ ), where $\operatorname{chars}[i]=b_{i}$, range $\left(W b_{i}\right)=[$ first $[i] . . f i r s t[i+1]-1]$ for $i \in[1 . . k]$, and range () refers to $\mathrm{BWT}_{T}$. Note that $\operatorname{range}(W)=[$ first $[1] . . f i r s t[k+1]-1]$, since it coincides with the concatenation of the intervals of the right extensions of $W$ in lexicographic order. If $W$ is not right-maximal, array chars in repr $(W)$ has length one. Given a data structure that supports rangeDistinct queries on $\mathrm{BWT}_{T}$, and given the $C$ array of $T$, there is an algorithm that converts $\operatorname{repr}(W)$ into the sequence $a_{1}, \ldots, a_{h}$ and into the corresponding sequence repr $\left(a_{1} W\right), \ldots, \operatorname{repr}\left(a_{h} W\right)$, in $O(d e)$ time and $O\left(\sigma^{2} \log n\right)$ bits of space in addition to the input and the output [3], where $d$ is the time taken by the rangeDistinct operation per element in its output, and $e$ is the number of distinct strings $a_{i} W b_{j}$ that occur in the circular version of $T$, where $i \in[1 . . h]$ and $j \in[1 . . k]$. We encapsulate this algorithm into a function that we call extendLeft.

If $a_{i} W$ is right-maximal, i.e. if array chars in $\operatorname{repr}\left(a_{i} W\right)$ has length greater than one, we push pair (repr $\left.\left(a_{i} W\right),|W|+1\right)$ onto a stack $S$. In the next iteration we pop the representation of a string from the stack and we repeat the process, until the stack itself becomes empty. This process is equivalent to following all the explicit Weiner links from the node $v$ of $\mathrm{ST}_{T}$ with $\ell(v)=W$, not necessarily in lexicographic order. Thus, running the algorithm from a stack initialized with $\operatorname{repr}(\varepsilon)$ is equivalent to performing a depth-first (but not necessarily a preorder) traversal of the suffix-link tree of $T$, which guarantees to enumerate all the right-maximal substrings of $T$. Every operation performed by the algorithm can be charged to a distinct node or Weiner link of $\mathrm{ST}_{T}$, thus the algorithm runs in $O(n d)$ time. The depth of the stack is $O(\log n)$ rather than $O(n)$, since at every iteration we push the pair (repr $\left.\left(a_{i} W\right),\left|a_{i} W\right|\right)$ with largest range $\left(a_{i} W\right)$ first. Every suffix-link tree level in the stack contains at most $\sigma$ pairs, and each pair takes at most $\sigma \log n$ bits of space, thus the total space used by the stack is $O\left(\sigma^{2} \log ^{2} n\right)$ bits. The following theorem follows from our assumption that $\sigma \in o(\sqrt{n} / \log n)$ :

Theorem 1 ([3]). Let $T \in[1 . . \sigma]^{n-1} \#$ be a string. Given a data structure that supports rangeDistinct queries on $\mathrm{BWT}_{T}$, we can enumerate all the right-maximal substrings $W$ of $T$, and for each of them we can return $|W|$, repr $(W)$, the sequence $a_{1}, a_{2}, \ldots, a_{h}$ of all characters in $\Sigma_{T}^{\ell}(W)$ (not necessarily sorted), and the sequence $\operatorname{repr}\left(a_{1} W\right), \ldots, \operatorname{repr}\left(a_{h} W\right)$, in $O(n d)$ time and in o(n) bits of space in addition to the input and the output, where $d$ is the time taken by the rangeDistinct operation per element in its output.

Theorem 1 does not specify the order in which the right-maximal substrings must be enumerated, nor the order in which the left extensions of a right-maximal substring must be returned. The algorithm we just described can be adapted to return all the maximal repeats of $T$, with the same bounds, by outputting a right-maximal string $W$ iff $\mid$ rangeDistinct $(\operatorname{sp}(W), \operatorname{ep}(W)) \mid>1$. A version of the same algorithm can also enumerate all the internal nodes of the generalized suffix tree of two string $T^{1}$ and $T^{2}$, using $\mathrm{BWT}_{T^{1}}$ and $\mathrm{BWT}_{T^{2}}$ : in this case, a string $W$ is represented as a quadruple $\operatorname{repr}^{\prime}(W)=\left(\operatorname{chars}_{1}\left[1 . . k_{1}\right]\right.$, first $t_{1}\left[1 . . k_{1}+\right.$ $1]$, chars $_{2}\left[1 . . k_{2}\right]$, first $_{2}\left[1 . . k_{2}+1\right]$ ), and we assume that first ${ }_{i}[1]=0$ iff $W$ does not occur in $T^{i}$. We call extendLeft' the function that maps repr ${ }^{\prime}(W)$ to the list of its left extensions repr ${ }^{\prime}\left(a_{i} W\right)$.

Theorem 2 ([3]). Let $T^{1} \in[1 . . \sigma]^{n_{1}-1} \#_{1}$ and $T^{2} \in[1 . . \sigma]^{n_{2}-1} \#_{2}$ be two strings. Given two data structures that support rangeDistinct queries on $\mathrm{BWT}_{T^{1}}$ and on $\mathrm{BWT}_{T^{2}}$, respectively, we can enumerate all the right-maximal substrings $W$ of $T=T^{1} T^{2}$, and for each of them we can return $|W|$, $\operatorname{repr}^{\prime}(W)$, the sequence $a_{1}, a_{2}, \ldots, a_{h}$ of all characters in $\Sigma_{T^{1} T^{2}}^{\ell}(W)$ (not necessarily sorted), and the sequence $\operatorname{repr}{ }^{\prime}\left(a_{1} W\right), \ldots, \operatorname{repr}^{\prime}\left(a_{h} W\right)$, in $O(n d)$ time and in $o(n)$ bits of space in addition to the input and the output, where $n=n_{1}+n_{2}$ and $d$ is the time taken by the rangeDistinct operation per element in its output.

For reasons of space, we assume throughout the paper that $d$ is the time per element in the output of a rangeDistinct data structure that is implicit from the context. We also replace $T^{i}$ by $i$ in subscripts, or we waive subscripts completely whenever they are clear from the context.

## 3 Kernels and complexity measures on $k$-mers

Given a string $T \in[1 . . \sigma]^{n-1} \#$ and a length $k>0$, let vector $\mathbf{T}_{k}=\left[1 . . \sigma^{k}\right]$ be such that $\mathbf{T}_{k}[W]=f_{T}(W)$ for every $W \in[1 . . \sigma]^{k}$. The $k$-mer complexity $\mathrm{C}_{k}(T)$ of string $T$ is the number of nonzero components of $\mathbf{T}_{k}$. The $k$-mer kernel of two strings $T^{1}$ and $T^{2}$ is $\kappa\left(\mathbf{T}_{k}^{1}, \mathbf{T}_{k}^{2}\right)$. Recall that Theorem 1 and 2 enumerate all nodes of a suffix tree in no specific order. In this section we describe algorithms to compute $\mathrm{C}_{k}(T)$ and $\kappa\left(\mathbf{T}_{k}^{1}, \mathbf{T}_{k}^{2}\right)$ in a way that does not depend on the order in which the nodes of a suffix tree are enumerated: we can thus implement such algorithms on top of Theorem 1 and 2. The main idea behind our approach is a telescoping strategy that works by adding and subtracting terms in a sum, as described below:

Theorem 3. Let $T \in[1 . . \sigma]^{n-1} \#$ be a string. Given an integer $k$ and a data structure that supports rangeDistinct queries on $\mathrm{BWT}_{T}$, we can compute $\mathrm{C}_{k}(T)$ in $O(n d)$ time and in o(n) bits of space in addition to the input.

Proof. A $k$-mer of $T$ can either be the label of a node of $\mathrm{ST}_{T}$, or it could end in the middle of an edge $(u, v)$ of ST. In the latter case, we assume that the $k$-mer is represented by its locus $v$, which might be a leaf. Let $\mathrm{C}_{k}(T)$ be initialized to $n-k$, i.e. to the number of leaves that correspond to suffixes of $T$ of length at least $k+1$. We enumerate the internal nodes of ST using Theorem 1, and every time we enumerate a node $v$ we proceed as follows: if $|\ell(v)|<k$ we leave $\mathrm{C}_{k}(T)$ unaltered, otherwise we increment $\mathrm{C}_{k}(T)$ by one and we decrement $\mathrm{C}_{k}(T)$ by the number of children of $v$ in ST , which is the length of array chars in repr $(\ell(v))$. In this way, every internal node $v$ of ST that is located at string depth at least $k$ and that is not the locus of a $k$-mer is both added to $\mathrm{C}_{k}(T)$ (when the algorithm visits $v$ ) and subtracted from $\mathrm{C}_{k}(T)$ (when the algorithm visits parent $(v))$. Leaves at depth at least $k+1$ that are not the locus of a $k$-mer are added by the initialization of $\mathrm{C}_{k}(T)$, and they are subtracted during the enumeration. Conversely, every locus $v$ of a $k$-mer of $T$ (including leaves) is just added to $\mathrm{C}_{k}(T)$, since $\mid \ell($ parent $(v)) \mid<k$.

We can apply the same telescoping strategy to compute $\kappa\left(\mathbf{T}_{k}^{1}, \mathbf{T}_{k}^{2}\right)$ :
Theorem 4. Let $T^{1} \in[1 . . \sigma]^{n_{1}-1} \#_{1}$ and $T^{2} \in[1 . . \sigma]^{n_{2}-1} \#_{2}$ be strings. Given an integer $k$ and two data structures that support rangeDistinct queries on $\mathrm{BWT}_{T^{1}}$ and on $\mathrm{BWT}_{T^{2}}$, respectively, we can compute $\kappa\left(\mathbf{T}_{k}^{1}, \mathbf{T}_{k}^{2}\right)$ in $O(n d)$ time and in o(n) bits of space in addition to the input, where $n=n_{1}+n_{2}$.

Proof. Recall that $\kappa\left(\mathbf{T}_{k}^{1}, \mathbf{T}_{k}^{2}\right)=N / \sqrt{D^{1} D^{2}}$, where $N=\sum_{W} \mathbf{T}_{k}^{1}[W] \mathbf{T}_{k}^{2}[W], D^{i}=\sum_{W} \mathbf{T}_{k}^{i}[W]^{2}$, and $W \in$ [1.. $\sigma]^{k}$. We initially set $N=0$ and $D^{i}=n_{i}-k$, since these are the contributions of all the leaves at depth at least $k+1$ in the generalized suffix tree of $T^{1}$ and $T^{2}$. Then, we enumerate every internal node $u$ of the generalized suffix tree, using Theorem 2 if $|\ell(u)|<k$ we keep all variables unchanged, otherwise we set $N$ to $N+f_{1}(\ell(u)) \cdot f_{2}(\ell(u))-\sum_{v} f_{1}(\ell(v)) \cdot f_{2}(\ell(v))$ and we set $D^{i}$ to $D^{i}+f_{i}(\ell(u))^{2}-\sum_{v} f_{i}(\ell(v))^{2}$, where $v$ ranges over all children of $u$ in the generalized suffix tree. Clearly $f_{i}(\ell(u))=\mathrm{first}_{i}\left[k_{i}+1\right]-\mathrm{first}_{i}[1]$ where $k_{i}$ is the size of $\operatorname{array~}_{c h a r s_{i}}$ in $\operatorname{repr}^{\prime}(\ell(u))$, and $f_{i}(\ell(v))=f_{i}\left(\ell(u) b_{j}\right)=\mathrm{first}_{i}[j+1]-\mathrm{first}_{i}[j]$ for some $j \in\left[1 . . k_{i}\right]$. In analogy to Lemma 3, the contribution of the loci of the distinct $k$-mers of $T^{1}$, of $T^{2}$, or of both, is added to the three temporary variables and never subtracted, while the contribution of every other node $u$ at depth at least $k$ in the generalized suffix tree is both added (when the algorithm visits $u$, or when $N$ and $D^{i}$ are initialized) and subtracted (when the algorithm visits parent (u)).

An even more specific notion of compositional complexity is $\mathrm{C}_{k, f}(T)$, the number of distinct $k$-mers that occur exactly $f$ times in $T$. In the $k$-mer profiling problem [6, 7] we are given a string $T$, an interval $\left[k_{1} . . k_{2}\right]$ of lengths and an interval $\left[f_{1} . . f_{2}\right]$ of frequencies, and we are asked to compute the matrix profile $\left[k_{1} . . k_{2}, f_{1} . . f_{2}\right]$ defined as follows: profile $[i, j]=\mathrm{C}_{i, j}(T)$ if $j<f_{2}$, and profile $[i, j]=\sum_{h \geq j} \mathrm{C}_{i, h}(T)$ if $j=f_{2}$. Note that column $j$ of profile can have nonzero cells only if $f_{j}$ is the frequency of some internal node of $\mathrm{ST}_{T}$. In practice profile is often computed by running a $k$-mer extraction algorithm $k_{2}-k_{1}+1$ times, and by scanning the output of all such runs (see e.g. [6] and references therein). The following lemma shows that we can compute profile in just one pass over the BWT of the input string, and in linear time in the size of profile:

Theorem 5. Let $T \in[1 . . \sigma]^{n-1} \#$ be a string. Given ranges $\left[k_{1} . . k_{2}\right]$ and $\left[f_{1} . . f_{2}\right]$, and given a data structure that supports rangeDistinct queries on $\mathrm{BWT}_{T}$, we can compute matrix profile $\left[k_{1} . . k_{2}, f_{1} . . f_{2}\right]$ in $O(n d+$ $\left.\left(k_{2}-k_{1}\right)\left(f_{2}-f_{1}\right)\right)$ time and in o(n) bits of space in addition to the input and the output.

Proof. We use Theorem 1 again. Assume that, for every internal node $u$ of $\mathrm{ST}_{T}$ with string depth at least $k_{1}$ and with frequency at least $f_{1}$, and for every $k \in\left[k_{1} . . \min \left\{|\ell(v)|, k_{2}\right\}\right]$, we increment profile $\left[k, \min \left\{f(u), f_{2}\right\}\right]$ by one and we decrement profile $\left[k, \min \left\{f(v), f_{2}\right\}\right]$ by one for every child $v$ of $u$ in ST such that $f(v) \geq f_{1}$. This would take $O\left(n^{2}\right)$ total updates to profile. However, we can perform all of these updates in batch, as follows: for every node $u$ of ST with $f(u) \geq f_{1}$ and with $|\ell(u)| \geq k_{1}$, we just increment profile $\left[\min \left\{|\ell(u)|, k_{2}\right\}, \min \left\{f(u), f_{2}\right\}\right]$ by one, and we just decrement profile $\left[\min \left\{|\ell(u)|, k_{2}\right\}, \min \left\{f(v), f_{2}\right\}\right]$ by one for every child $v$ of $u$ in ST such that $f(v) \geq f_{1}$. After having traversed all the internal nodes of ST, we scan profile as follows: for every $j \in\left[f_{1} . . f_{2}\right]$, we traverse all values of $i$ in the decreasing order $k_{2}-1, \ldots, k_{1}$, and we set profile $[i, j]=\operatorname{profile}[i, j]+\operatorname{profile}[i+1, j]$. If $f_{1}=1$, at the end of this process the first column of profile contains negative numbers, since Theorem 1 does not enumerate the leaves of ST. Thus, before returning, we add to profile $[i, 1]$ the number of leaves with string depth at least $k_{i}+1$, i.e. value $n-k_{i}$, for all $i \in\left[k_{1} . . k_{2}\right]$.

A similar algorithm allows computing $\kappa\left(\mathbf{T}_{k}^{1}, \mathbf{T}_{k}^{2}\right)$ for all $k$ in a user-specified range $\left[k_{1} . . k_{2}\right]$ in $O\left(n d+k_{2}-k_{1}\right)$ time. Matrix profile can be used to determine a range of values of $k$ to be used in $k$-mer kernels. The smallest number in this range is typically the value of $k$ that maximizes the number of distinct $k$-mers that occur at least twice in $T$ [15]. The largest number in the range is typically determined using some measure of expectation: we cover this computation in Section 5

A related notion of compositional complexity is the $k$-th order empirical entropy of $T$, defined as $\mathrm{H}_{k}(T)=$ $(1 /|T|) \cdot \sum_{W} \sum_{a \in \Sigma^{r}(W)} f_{T}(W a) \cdot \log \left(f_{T}(W) / f_{T}(W a)\right)$, where $W$ ranges over all strings in $[1 . . \sigma]^{k}$. Clearly only the internal nodes of $\mathrm{ST}_{T}$ contribute to some $\mathrm{H}_{k}(T)$ [9] thus our methods allow computing $\mathrm{H}_{k}(T)$ for a user-specified range of lengths $\left[k_{1} . . k_{2}\right]$ in $O(n d)$ time, using just one pass over $\mathrm{BWT}_{T}$.

## 4 Kernels and complexity measures on all substrings

Given a string $T \in[1 . . \sigma]^{n-1} \#$, consider the infinite-dimensional vector $\mathbf{T}_{\infty}$ indexed by all distinct substrings $W \in[1 . . \sigma]^{+}$, such that $\mathbf{T}_{\infty}[W]=f_{T}(W)$. The substring complexity $\mathrm{C}_{\infty}(T)$ of $T$ is the number of nonzero components of $\mathbf{T}_{\infty}$. The substring kernel of two strings $T^{1}$ and $T^{2}$ is the cosine of composition vectors $\mathbf{T}_{\infty}^{1}$ and $\mathbf{T}_{\infty}^{2}$. Computing substring complexity and substring kernel amounts to applying the same telescoping strategy described in Theorem 3 and [4] but with different contributions:

Corollary 1. Let $T \in[1 . . \sigma]^{n-1} \#$ be a string. Given a data structure that supports rangeDistinct queries on $\mathrm{BWT}_{T}$, we can compute $\mathrm{C}_{\infty}(T)$ in $O(n d)$ time and in $o(n)$ bits of space in addition to the input.

Proof. The substring complexity of $T$ coincides with the number of characters in [1.. $\sigma$ ] that occur on all edges of $\mathrm{ST}_{T}$. We can thus proceed as in Lemma [3, initializing $\mathrm{C}_{\infty}(T)$ to $(n-1) n / 2$, or equivalently to the sum of the lengths of all suffixes of $T[1 . . n-1]$. Whenever we visit a node $v$ of ST , we add to $\mathrm{C}_{\infty}(T)$ the quantity $|\ell(v)|$, and we subtract from $\mathrm{C}_{\infty}(T)$ the quantity $|\ell(v)| \cdot|\operatorname{children}(v)|$. The net effect of all such operations coincides with summing the lengths of all edges of ST, discarding all occurrences of character \#. Note that $|\ell(u)|$ is provided by Theorem [1, and $|\operatorname{children}(v)|$ is the size of array chars in $\operatorname{repr}(\ell(v))$.

Corollary 2. Let $T^{1} \in[1 . . \sigma]^{n_{1}-1} \#_{1}$ and $T^{2} \in[1 . . \sigma]^{n_{2}-1} \#_{2}$ be strings. Given data structures that support rangeDistinct queries on $\mathrm{BWT}_{T^{1}}$ and on $\mathrm{BWT}_{T^{2}}$, respectively, we can compute $\kappa\left(\mathbf{T}_{\infty}^{1}, \mathbf{T}_{\infty}^{2}\right)$ in $O(n d)$ time and in o( $n$ ) bits of space in addition to the input, where $n=n_{1}+n_{2}$.

Proof. We proceed as in Theorem 4 setting again $N=0$ and $D^{i}=\left(n_{i}-1\right) n_{i} / 2$ at the beginning of the algorithm. When we visit a node $u$ of the generalized suffix tree of $T^{1}$ and $T^{2}$, we set $N$ to $N+|\ell(u)|$. $\left(f_{1}(\ell(u)) f_{2}(\ell(u))-\sum_{v} f_{1}(\ell(v)) f_{2}(\ell(v))\right)$ and we set $D^{i}$ to $D^{i}+|\ell(u)| \cdot\left(f_{i}(\ell(u))^{2}-\sum_{v} f_{i}(\ell(v))^{2}\right)$, where $v$ ranges over all children of $u$ in the generalized suffix tree.

In a substring kernel it is common to weight a substring $W$ by a user-specified function of its length: typical choices are $\epsilon^{|W|}$ for a given constant $\epsilon$, or indicators that select only substrings within a specific range of lengths [16. We denote by $\mathbf{T}_{\infty, g}^{i}$ a weighted version of the infinite-dimensional vector $\mathbf{T}_{\infty}^{i}$ in which $\mathbf{T}_{\infty}^{i}[W]=g(|W|) \cdot f_{T^{i}}(W)$ and where $g$ is any user-specified function. We assume that the number of bits required to represent the output of $g$ with sufficient precision is $O(\log n)$. It is easy to adapt Corollary 2 to support this type of composition vector:
Corollary 3. Let $T^{1} \in[1 . . \sigma]^{n_{1}-1} \#_{1}$ and $T^{2} \in[1 . . \sigma]^{n_{2}-1} \#_{2}$ be strings. Given a function $g(k)$ that can be evaluated in constant time, and given data structures that support rangeDistinct queries on $\mathrm{BWT}_{T^{1}}$ and on $\mathrm{BWT}_{T^{2}}$, respectively, we can compute $\kappa\left(\mathbf{T}_{\infty, g}^{1}, \mathbf{T}_{\infty, g}^{2}\right)$ in $O(n d)$ time and in o(n) bits of space in addition to the input, where $n=n_{1}+n_{2}$.

Proof. We modify Corollary 2 as follows. Assume that we are processing an internal node $v$ of the generalized suffix tree, let $\ell(v)=W$, and assume that we have computed repr' $(a W)$ for all the left extensions $a W$ of $W$. In addition to pushing $\operatorname{repr}^{\prime}(a W)$ onto the stack, we also push value prefixSum $(a W)=\sum_{i=1}^{|W|+1} g(i)^{2}$ with it, where prefixSum $(a W)=\operatorname{prefixSum}(W)+g(|W|+1)^{2}$. When we pop repr${ }^{\prime}(a W)$, we compute its contributions to $N$ and $D^{i}$ as described in Corollary 2, but replacing $|a W|$ by prefixSum $(a W)$. We initialize $D^{i}$ to $\sum_{j=1}^{n_{i}-1} g(j)^{2}$.

Corollary 3 can clearly support distinct weight functions for $T^{1}$ and $T^{2}$. For some functions, like $\epsilon^{|W|}$, prefix sums can be computed in closed form [16], thus there is no need to push prefixSum values on the stack. Another frequent weighting scheme for a string $W$ associates a score $q(c)$ to every character $c$ of $W$, and it weights $W$ by e.g. $q(W)=\prod_{i=1}^{|W|} q(W[i])$. In this case we could just push prefixSum $(V)=$ $\sum_{i=1}^{|V|} \prod_{j=1}^{i} q(V[j])^{2}$ onto the stack, where $V=a W$ and prefixSum $(V)=q(a)^{2} \cdot(1+\operatorname{prefixSum}(W))$. A similar weighting scheme can be used for $k$-mers as well. Let $\mathbf{T}_{k, q}$ be a version of $\mathbf{T}_{k}$ such that $\mathbf{T}_{k, q}[W]=$ $f_{T}(W)-(|T|-|W|) q(W)$ for every $W \in[1 . . \sigma]^{k}$, and consider the following distances defined in [13]:

$$
\begin{aligned}
D_{2}^{s}\left(\mathbf{T}_{k, q}^{1}, \mathbf{T}_{k, q}^{2}\right) & =\sum_{W} \mathbf{T}_{k, q}^{1}[W] \mathbf{T}_{k, q}^{2}[W] / \sqrt{\left(\mathbf{T}_{k, q}^{1}[W]\right)^{2}+\left(\mathbf{T}_{k, q}^{2}[W]\right)^{2}} \\
D_{2}^{*}\left(\mathbf{T}_{k, q}^{1}, \mathbf{T}_{k, q}^{2}\right) & =\sum_{W} \mathbf{T}_{k, q}^{1}[W] \mathbf{T}_{k, q}^{2}[W] /\left(\sqrt{\left(n_{1}-k\right)\left(n_{2}-k\right)} \cdot q(W)\right)
\end{aligned}
$$

where $W$ ranges over all strings in $[1 . . \sigma]^{k}$. We can compute such distances using just a minor modification to Theorem 4

Corollary 4. Let $T^{1} \in[1 . . \sigma]^{n_{1}-1} \#_{1}$ and $T^{2} \in[1 . . \sigma]^{n_{2}-1} \#_{2}$ be strings. Given an integer $k$ and data structures that support rangeDistinct queries on $\mathrm{BWT}_{T^{1}}$ and on $\mathrm{BWT}_{T^{2}}$, respectively, we can compute $D_{2}^{s}\left(\mathbf{T}_{k, p}^{1}, \mathbf{T}_{k, p}^{2}\right)$ and $D_{2}^{*}\left(\mathbf{T}_{k, p}^{1}, \mathbf{T}_{k, p}^{2}\right)$ in $O(n d)$ time and in $\lambda \log \sigma+o(n)$ bits of space in addition to the input, where $n=n_{1}+n_{2}$ and $\lambda$ is the length of the longest repeat in $T^{1} T^{2}$.
Proof. We proceed as in Theorem 4. pushing on the stack value $q(W, k)=\prod_{j=1}^{k} q(W[j])$ in addition to $\operatorname{repr}^{\prime}(W)$, and maintaining a separate stack of characters to represent the string we are processing during the depth-first traversal of the generalized suffix-link tree. We set $q(a W, k)=q(a) \cdot q(W, k) / q(b)$, where $b$ is the $k$ th character from the top of the character stack when we are processing $W$.

An orthogonal way to measure the similarity between $T^{1}$ and $T^{2}$ consists in comparing the repertoire of all strings that do not appear in $T^{1}$ and in $T^{2}$. Given a string $T$ and two frequency thresholds $\tau_{1}<\tau_{2}$, a string $W$ is a minimal rare word of $T$ if $\tau_{1} \leq f_{T}(W)<\tau_{2}$ and if $f_{T}(V) \geq \tau_{2}$ for every proper substring $V$ of $W$. Setting $\tau_{1}=0$ and $\tau_{2}=1$ gives the well-known minimal absent words (see e.g. [10, [5] and references therein), whose total number can be $\Theta(\sigma n)$ 8. Setting $\tau_{1}=1$ and $\tau_{2}=2$ gives the so-called shortest unique substrings (see e.g. [11] and references therein), whose total number is $O(n)$, like the number of strings obtained by any other setting of $\tau_{1} \geq 1$. In what follows we focus on minimal absent words, but our algorithms can be generalized to other settings of the thresholds.

To decide whether $a W b$ is a minimal absent word of $T$, where $a$ and $b$ are characters, it clearly suffices to check whether $f_{T}(a W b)=0$ and whether both $f_{T}(a W) \geq 1$ and $f_{T}(W b) \geq 1$. It is well known that only a maximal repeat of $T$ can be the infix $W$ of a minimal absent word $a W b$, and this applies to any setting of $\tau_{1}$ and $\tau_{2}$. To enumerate all the minimal absent words, for example to count their total number $C_{-}(T)$, we can thus iterate over all nodes of $\mathrm{ST}_{T}$ associated with maximal repeats, as described below:

Theorem 6. Let $T \in[1 . . \sigma]^{n-1} \#$ be a string. Given a data structure that supports rangeDistinct queries on $\mathrm{BWT}_{T}$, we can compute $\mathrm{C}_{-}(T)$ in $O(n d)$ time and in $o(n)$ bits of space in addition to the input.

Proof. For clarity, we first describe how to enumerate all the distinct minimal absent words of $T$ : we specialize this algorithm to counting at the end of the proof. We use Theorem 1 to enumerate all nodes $v$ of $\mathrm{ST}_{T}$ associated with maximal repeats, as described in Section [2.2, Let $\left\{a_{1}, \ldots, a_{h}\right\}$ be the set of distinct left extensions of string $\ell(v)$ in $T$ returned by operation extendLeft $(\operatorname{repr}(v))$, let extensions $[1 . . \sigma+1,0 . . \sigma]$ be a boolean matrix initialized to all zeros, and let leftExtensions $[1 . . \sigma+1]$ be an array initialized to all zeros. Let $h^{\prime}$ be a pointer initialized to one. Operation extendLeft allows following all the Weiner links from $v$, not necessarily in lexicographic order: for every string $a_{i} \ell(v)$ obtained in this way, we set leftExtensions $\left[h^{\prime}\right]=a_{i}$, we enumerate its right extensions $\left\{c_{1}, \ldots, c_{k^{\prime}}\right\}$ using array chars of repr $\left(a_{i} \ell(v)\right)$, we set extensions $\left[h^{\prime}, c_{j}\right]=1$ for all $j \in\left[1 . . k^{\prime}\right]$, and we finally increment $h^{\prime}$ by one. Note that only the columns of extensions that correspond to the right extensions of $\ell(v)$ are updated by this procedure. Then, we enumerate all the right extensions $\left\{b_{1}, \ldots, b_{k}\right\}$ of $\ell(v)$ using array chars of repr $(\ell(v))$, and for every such extension $b_{j}$ we report all pairs $\left(a_{i}, b_{j}\right)$ such that $a_{i}=\operatorname{chars}[x], x \in\left[1 . . h^{\prime}\right]$, and extensions $\left[x, b_{j}\right]=0$. This process takes time proportional to the number of Weiner links from $v$, plus the number of children of $v$, plus the number of Weiner links from $v$ multiplied by $\sigma$. When applied to all nodes of ST , this takes in total $O(n \sigma)$ time, which is optimal in the size of the output. The matrices and vectors used by this process can be reset to all zeros after processing each node: the total time spent in such reinitializations in $O(n)$.

If we just need $C_{-}(T)$, rather than storing the temporary matrices extensions and leftExtensions, we store just a number area which we initialize to $h k$ before processing node $v$. Whenever we observe a right extension $c_{j}$ of a string $a_{i} \ell(v)$, we decrease area by one. Before moving to the next node, we increment $C_{-}(T)$ by area.

Let $\mathbf{T}_{-}$be the infinite-dimensional vector indexed by all distinct substrings $W \in[1 . . \sigma]^{+}$, such that $\mathbf{T}_{-}[W]=1$ iff $W$ is a minimal absent word of $T$. Theorem 6 can be adapted to compute the Jaccard distance between the composition vectors of two strings:

Corollary 5. Let $T^{1} \in[1 . . \sigma]^{n_{1}-1} \#_{1}$ and $T^{2} \in[1 . . \sigma]^{n_{2}-1} \#_{2}$ be strings. Given data structures that support rangeDistinct queries on $\mathrm{BWT}_{T^{1}}$ and on $\mathrm{BWT}_{T^{2}}$, respectively, we can compute $J\left(\mathbf{T}_{-}^{1}, \mathbf{T}_{-}^{2}\right)$ in $O(n d)$ time and in $o(n)$ bits of space in addition to the input, where $n=n_{1}+n_{2}$.

Proof. We apply the strategy of Theorem 6 to the internal nodes of the generalized suffix tree of $T^{1}$ and $T^{2}$ whose label is a maximal repeat of $T^{1}$ and a maximal repeat of $T^{2}$ : such strings are clearly maximal repeats of $T^{1} T^{2}$ as well. We enumerate such nodes as described in Section 2.2 We keep a global variable intersection and a bitvector sharedRight[1.. $\sigma]$. For every node $v$ that corresponds to a maximal repeat of $T^{1}$ and of $T^{2}$, we merge the sorted arrays chars $s_{1}$ and $\operatorname{chars}_{2}$ of $\operatorname{repr}^{\prime}(\ell(v))$, we set sharedRight $[c]=1$ for every character $c$ that belongs to the intersection of the two arrays, and we cumulate in a variable $k^{\prime}$ the number of ones in sharedRight. Then, we scan every left extension $a_{i}$ provided by extendLeft', we determine in constant time whether it occurs in both $T^{1}$ and $T^{2}$, and if so we increment a variable $h^{\prime}$ by one. Finally, we initialize a variable area to $h^{\prime} k^{\prime}$, and we process again every left extension $a_{i}$ provided by extendLeft': if $a_{i} \ell(v)$ occurs in both $T^{1}$ and $T^{2}$, we compute the union of arrays chars ${ }_{1}$ and chars ${ }_{2}$ of $\operatorname{repr}^{\prime}\left(a_{i} \ell(v)\right)$, and for every character $c$ in the union such that sharedRight $[c]=1$, we decrement area by one. At the end of this process, we add area to the global variable intersection. To compute $\left\|\mathbf{T}_{-}^{1} \vee \mathbf{T}_{-}^{2}\right\|$ we apply Theorem 6 to $T^{1}$ and $T^{2}$ separately.

It is easy to extend Corollary 5 to compute $\kappa\left(\mathbf{T}_{-}^{1}, \mathbf{T}_{-}^{2}\right)$, as well as to support weighting schemes based on the length and on the characters of minimal absent words.

## 5 Markovian corrections

In some applications it is desirable to assign to component $W \in[1 . . \sigma]^{k}$ of composition vector $\mathbf{T}_{\infty}$ an estimate of the statistical significance of observing $f_{T}(W)$ occurrences of $W$ in $T$ : intuitively, strings whose frequency departs from its expected value are more likely to carry "information", and they should be weighted more [12]. Assume that $T$ is generated by a Markov random process of order $k-2$ or smaller, that produces strings on alphabet $[1 . . \sigma]$ according to a probability distribution $\mathbb{P}$. It is well known that the probability of observing $W$ in a string generated by such random process is $\mathbb{P}(W)=\mathbb{P}(W[1 . . k-1]) \cdot \mathbb{P}(W[2 . . k]) / \mathbb{P}(W[2 . . k-1])$. We can estimate $\mathbb{P}(W)$ using the empirical probability $p_{T}(W)$, obtaining the following approximation for $\mathbb{P}(W)$ : $\tilde{p}_{T}(W)=p_{T}(W[1 . . k-1]) \cdot p_{T}(W[2 . . k]) / p_{T}(W[2 . . k-1])$ if $p_{T}(W[2 . . k-1]) \neq 0$, and $\tilde{p}_{T}(W)=0$ otherwise. We can thus estimate the significance of the event that substring $W$ has empirical probability $p_{T}(W)$ in string $T$ using the following score: $z_{T}(W)=\left(p_{T}(W)-\tilde{p}_{T}(W)\right) / \tilde{p}_{T}(W)$ if $\tilde{p}_{T}(W) \neq 0$, and $z_{T}(W)=0$ if $\tilde{p}_{T}(W)=0$ [12]. After elementary manipulations [2], $z_{T}(W)$ becomes:

$$
\begin{aligned}
z_{T}(W) & =g(n, k) \cdot \frac{f_{T}(W) \cdot f_{T}(W[2 . . k-1])}{f_{T}(W[1 . . k-1]) \cdot f_{T}(W[2 . . k])}-1 \\
g(x, y) & =(x-y+2)^{2} /(x-y+1)(x-y+3)
\end{aligned}
$$

Since $g(x, y) \in[1 . .1 .125]$, we temporarily assume $g(x, y)=1$ in what follows, removing this assumption later.
Let $\mathbf{T}_{z}$ be a version of the infinite-dimensional vector $\mathbf{T}_{\infty}$ in which $\mathbf{T}_{z}[W]=z_{T}(W)$. Among all strings that occur in $T$, only strings $a W b$ such that $a$ and $b$ are characters in $[0 . . \sigma]$ and such that $W$ is a maximal repeat of $T$ can have $\mathbf{T}_{z}[a W b] \neq 0$. Similarly, among all strings that do not occur in $T$, only the minimal absent words of $T$ have a nonzero component in $\mathbf{T}_{z}$ : specifically, $\mathbf{T}_{z}[a W b]=-1$ for all minimal absent words $a W b$ of $T$, where $a$ and $b$ are characters in $[0 . . \sigma][2]$. Given two strings $T^{1}$ and $T^{2}$, we can thus compute $\kappa\left(\mathbf{T}_{z}^{1}, \mathbf{T}_{z}^{2}\right)$ using the same strategy as in Corollary 5

Theorem 7. Let $T^{1} \in[1 . . \sigma]^{n_{1}-1} \#_{1}$ and $T^{2} \in[1 . . \sigma]^{n_{2}-1} \#_{2}$ be strings. Given data structures that support rangeDistinct queries on $\mathrm{BWT}_{T^{1}}$ and on $\mathrm{BWT}_{T^{2}}$, respectively, and assuming $g(x, y)=1$ for all settings of $x$ and $y$, we can compute $\kappa\left(\mathbf{T}_{z}^{1}, \mathbf{T}_{z}^{2}\right)$ in $O(n d)$ time and in $o(n)$ bits of space in addition to the input, where $n=n_{1}+n_{2}$.

Proof. We focus here on computing component $N$ of $\kappa\left(\mathbf{T}_{z}^{1}, \mathbf{T}_{z}^{2}\right)$ : computing $D^{i}$ follows a similar algorithm on $\mathrm{BWT}_{T^{i}}$. We keep again a bitvector sharedRight[1.. $\left.\sigma\right]$, and we enumerate all the internal nodes of the generalized suffix tree of $T^{1}$ and $T^{2}$ whose label is a maximal repeat of $T^{1} T^{2}$, as described in Section 2.2, For every such node $v$, we merge the sorted arrays chars ${ }_{1}$ and $\operatorname{chars}_{2}$ of $\operatorname{repr}^{\prime}(\ell(v))$, we set sharedRight $[c]=1$ for every character $c$ that belongs to the intersection of the two arrays, and we cumulate in a variable $k^{\prime}$ the number of ones in sharedRight. Then, we scan every left extension $a_{i}$ provided by extendLeft', we determine in constant time whether it occurs in both $T^{1}$ and $T^{2}$, and if so we increment a variable $h^{\prime}$ by one. Finally, we initialize a variable area to $h^{\prime} k^{\prime}$, and we process again every left extension $a_{i}$ provided by extendLeft'. If $a_{i} \ell(v)$ occurs in both $T^{1}$ and $T^{2}$, we merge arrays chars ${ }_{1}$ and chars of repr $^{\prime}\left(a_{i} \ell(v)\right)$ : for every character $b$ in the intersection of chars $s_{1}$ and $\operatorname{chars}_{2}$, we add to $N$ value $z_{1}\left(a_{i} \ell(v) b\right) \cdot z_{2}\left(a_{i} \ell(v) b\right)$, retrieving the corresponding frequencies from $\operatorname{repr}^{\prime}\left(a_{i} \ell(v)\right)$ and from $\operatorname{repr}^{\prime}(\ell(v))$, and we decrement area by one. For every character $b$ that occurs only in chars ${ }_{1}$, we test whether sharedRight $[b]=1$ : if so, $a_{i} W b$ is a minimal absent word of $T^{2}$ that occurs in $T^{1}$, thus we decrement area by one and we add to $N$ value $-z_{1}\left(a_{i} \ell(v) b\right)$. We proceed symmetrically if $b$ occurs only in chars $s_{2}$. At the end of this process, area counts the number of minimal absent words with infix $\ell(v)$ that are shared by $T^{1}$ and $T^{2}$ : thus, we add area to $N$.

It is easy to remove the assumption that $g(x, y)$ is always equal to one. There are only two differences from the previous case. First, the score of the substrings $W$ of $T^{i}$ that have a maximal repeat of $T^{i}$ as an infix changes, but $g\left(n_{i},|W|\right)$ can be immediately computed from $|W|$, which is included in both repr $(W)$ and $\operatorname{repr}^{\prime}(W)$. Second, the score of all substrings $W$ of $T^{i}$ that do not have a maximal repeat as an infix changes from zero to $g\left(n_{i},|W|\right)-1$ : we can take into account all such contributions by pushing prefix-sums
to the stack, as in Corollary 3. For example, to compute component $N$ of $\kappa\left(\mathbf{T}_{z}^{1}, \mathbf{T}_{z}^{2}\right)$, we can first assume that all substring $W$ that occur both in $T^{1}$ and in $T^{2}$ have score $g\left(n_{i},|W|\right)-1$, by pushing to the stack the prefix-sums described in 2] and by enumerating only nodes $v$ of the generalized suffix tree of $T^{1}$ and $T^{2}$ such that $\ell(v)$ occurs both in $T^{1}$ and in $T^{2}$. Then, we can run the algorithm in Theorem 77 subtracting quantity $\left(g\left(n_{1},|W|+2\right)-1\right) \cdot\left(g\left(n_{2},|W|+2\right)-1\right)$ from the contribution to $N$ of every string $a_{i} W b$ that occurs both in $T^{1}$ and in $T^{2}$.

Finally, recall that in Section 3 we mentioned the problem of determining an upper bound on the values of $k$ to be used in $k$-mer kernels. Let $\mathbf{T}_{k}$ be the composition vector indexed by all strings in $[1 . . \sigma]^{k}$ such that $\mathbf{T}_{k}[W]=p_{T}(W)$, and let $\tilde{\mathbf{T}}_{k}$ be a similar composition vector with $\tilde{\mathbf{T}}_{k}[W]=\tilde{p}_{T}(W)$, where $\tilde{p}_{T}(W)$ is defined as in the beginning of this section. It makes sense to disregard values of $k$ for which $\mathbf{T}_{k}$ and $\tilde{\mathbf{T}}_{k}$ are very similar, and more formally whose Kullback-Leibler divergence $\mathrm{KL}\left(\mathbf{T}_{k}, \tilde{\mathbf{T}}_{k}\right)=\sum_{W} \mathbf{T}_{k}[W] \cdot\left(\log \left(\mathbf{T}_{k}[W]\right)-\right.$ $\left.\log \left(\tilde{\mathbf{T}}_{k}[W]\right)\right)$ is small, where $W$ ranges over all strings in $[1 . . \sigma]^{k}$. Thus, we could use as an upper bound on $k$ the minimum value $k^{*}$ such that $\sum_{k^{\prime}=k^{*}}^{\infty} \mathrm{KL}\left(\mathbf{T}_{k^{\prime}}, \tilde{\mathbf{T}}_{k^{\prime}}\right)<\tau$ for some user-specified threshold $\tau$ [15]. Note again that only strings $a W b$ such that $a$ and $b$ are characters in $[0 . . \sigma]$ and $W$ is a maximal repeat of $T$ contribute to $\mathrm{KL}\left(\mathbf{T}_{|W|+2}, \tilde{\mathbf{T}}_{|W|+2}\right)$. We can thus adapt Theorem 7 to compute the KL divergence for a user-specified range of lengths $\left[k_{1} . . k_{2}\right]$, using just one pass over $\mathrm{BWT}_{T}$, in $O(n d)$ time and in $o(n)$ bits of space in addition to the input and the output. The same approach can be used to compute the KL-divergence kernel $\kappa\left(\mathbf{T}_{K L}^{1}, \mathbf{T}_{K L}^{2}\right)$, where $\mathbf{T}_{K L}^{i}[W]=\mathrm{KL}_{T^{i}}(W)$ and $\mathrm{KL}_{T^{i}}(W)=\sum_{a, b \in \Sigma} p_{T^{i}}(a W b) \cdot\left(\log \left(p_{T^{i}}(a W b)\right)-\log \left(\tilde{p}_{T^{i}}(a W b)\right)\right)$.

## References

[1] Alberto Apostolico. Maximal words in sequence comparisons based on subword composition. In Algorithms and Applications, pages 34-44. Springer, 2010.
[2] Alberto Apostolico and Olgert Denas. Fast algorithms for computing sequence distances by exhaustive substring composition. Algorithms for Molecular Biology, 3(1):13, 2008.
[3] Djamal Belazzougui. Linear time construction of compressed text indices in compact space. In Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014, pages 148-193, 2014.
[4] Djamal Belazzougui, Gonzalo Navarro, and Daniel Valenzuela. Improved compressed indexes for fulltext document retrieval. Journal of Discrete Algorithms, 18:3-13, January 2013.
[5] Supaporn Chairungsee and Maxime Crochemore. Using minimal absent words to build phylogeny. Theoretical Computer Science, 450:109-116, 2012.
[6] Rayan Chikhi and Paul Medvedev. Informed and automated $k$-mer size selection for genome assembly. Bioinformatics, 30(1):31-37, 2014.
[7] Benny Chor, David Horn, Nick Goldman, Yaron Levy, Tim Massingham, et al. Genomic DNA $k$-mer spectra: models and modalities. Genome Biology, 10(10):R108, 2009.
[8] Maxime Crochemore, Filippo Mignosi, and Antonio Restivo. Automata and forbidden words. Information Processing Letters, 67(3):111-117, 1998.
[9] Simon Gog. Compressed suffix trees: Design, construction, and applications. PhD thesis, University of Ulm, Germany, 2011.
[10] Julia Herold, Stefan Kurtz, and Robert Giegerich. Efficient computation of absent words in genomic sequences. BMC Bioinformatics, 9(1):167, 2008.
[11] Atalay Mert Ileri and Bojian Xu. Shortest unique substring query revisited. In Combinatorial Pattern Matching, pages 172-181, 2014.
[12] Ji Qi, Bin Wang, and Bai-Iin Hao. Whole proteome prokaryote phylogeny without sequence alignment: a $k$-string composition approach. Journal of Molecular Evolution, 58(1):1-11, 2004.
[13] Gesine Reinert, David Chew, Fengzhu Sun, and Michael S Waterman. Alignment-free sequence comparison (I): statistics and power. Journal of Computational Biology, 16(12):1615-1634, 2009.
[14] John Shawe-Taylor and Nello Cristianini. Kernel methods for pattern analysis. Cambridge university press, 2004.
[15] Gregory E Sims, Se-Ran Jun, Guohong A Wu, and Sung-Hou Kim. Alignment-free genome comparison with feature frequency profiles (FFP) and optimal resolutions. Proceedings of the National Academy of Sciences, 106(8):2677-2682, 2009.
[16] Alex J. Smola and S.v.n. Vishwanathan. Fast kernels for string and tree matching. In S. Becker, S. Thrun, and K. Obermayer, editors, Advances in Neural Information Processing Systems 15, pages 585-592. MIT Press, 2003.


[^0]:    *This work was partially supported by Academy of Finland under grant 250345 (Center of Excellence in Cancer Genetics Research).

