# Fast Alignment-Free Similarity Estimation By Tensor Sketching 

Amir Joudaki ${ }^{1,2,3}$, Gunnar Rätsch ${ }^{1,2,3, *}$, and André Kahles ${ }^{1,2,3, *}$<br>${ }^{1}$ Department of Computer Science, ETH Zurich, Zurich 8092, Switzerland<br>\{ajoudaki,raetsch, andre.kahles\}@inf.ethz.ch<br>2 University Hospital Zurich, Biomedical Informatics Research, Zurich 8091, Switzerland<br>${ }^{3}$ SIB Swiss Institute of Bioinformatics, Lausanne 1015, Switzerland<br>* to whom correspondence should be addressed


#### Abstract

The sharp increase in next-generation sequencing technologies' capacity has created a demand for algorithms capable of quickly searching a large corpus of biological sequences. The complexity of biological variability and the magnitude of existing data sets have impeded finding algorithms with guaranteed accuracy that efficiently run in practice. Our main contribution is the Tensor Sketch method that efficiently and accurately estimates edit distances. In our experiments, Tensor Sketch had 0.88 Spearman's rank correlation with the exact edit distance, almost doubling the 0.466 correlation of the closest competitor while running 8.8 times faster. Finally, all sketches can be updated dynamically if the input is a sequence stream, making it appealing for large-scale applications where data cannot fit into memory. Conceptually, our approach has three steps: 1) represent sequences as tensors over their sub-sequences, 2) apply tensor sketching that preserves tensor inner products, 3 ) implicitly compute the sketch. The sub-sequences, which are not necessarily contiguous pieces of the sequence, allow us to outperform $k$-mer-based methods, such as min-hash sketching over a set of $k$-mers. Typically, the number of sub-sequences grows exponentially with the sub-sequence length, introducing both memory and time overheads. We directly address this problem in steps 2 and 3 of our method. While the sketching of rank-1 or super-symmetric tensors is known to admit efficient sketching, the sub-sequence tensor does not satisfy either of these properties. Hence, we propose a new sketching scheme that completely avoids the need for constructing the ambient space. Our tensor-sketching technique's main advantages are three-fold: 1) Tensor Sketch has higher accuracy than any of the other assessed sketching methods used in practice. 2) All sketches can be computed in a streaming fashion, leading to significant time and memory savings when there is overlap between input sequences. 3) It is straightforward to extend tensor sketching to different settings leading to efficient methods for related sequence analysis tasks. We view tensor sketching as a framework to tackle a wide range of relevant bioinformatics problems, and we are confident that it can bring significant improvements for applications based on edit distance estimation.


Keywords: sequence distance metrics, tensor sketching, edit distance estimation

## 1 Introduction

The emergence of next-generation sequencing technologies and a dramatic decrease in cost have led to an exponential increase in biological sequence data, frequently stored in exceptionally large databases. While alignment scores are considered the gold-standard of sequence distance in many bioinformatics applications, the growing number of sequences to be analyzed poses serious challenges to exact distance computation via alignment. This problem has motivated research on timeand space-efficient alignment-free methods that try to estimate rather than compute sequence similarity. Especially applications relying on comparing millions of sequences have turned towards choosing approximate methods over exact alignments [2]. Out of many possibilities, we have selected the task of phylogeny reconstruction to further motivate distance estimation. In this task, the estimated distances between sequences are used to reconstruct the evolutionary history in form of a phylogenetic tree. Instead of using exact alignment scores, many alignment-free methods instead rely on $k$-mer statistics as a proxy. (Where the term $k$-mer refers to all length $k$ substrings of a given string.) The multiplicity, frequency, mode, and reappearance statistics of $k$-mers have been utilised to directly estimate evolutionary distance [17|16,1]. Other approaches include variable lengths matches in their statistics [18].

To break the stringent dependency of adjacent letters in a $k$-mer, spaced $k$-mers have introduced "match" and "ignore" positions. For example, if a match-pattern " $\underline{11011}$ " is used, then " $\underline{T} G \underline{A C}$ "
versus "CTTAC" constitutes a match. Spaced $k$-mers have been shown to improve mapping sensitivity, the accuracy of phylogenies, and the performance of sequence classification [5, 7/13 15]. Analogously, substring-based methods can be relaxed to allow for some mismatches [6]. Since the quality of the estimations will greatly depend on the selected match-pattern, several works have focused on optimizing these patterns 311810. However, finding combinatorially optimal patterns becomes intractable as the number of ignore-positions increases. Furthermore, any optimization is prone to be task-dependent, which would require the optimization to be repeated for each task separately.

More recently, hashing based methods have become increasingly popular. MinHash sketch methods have been primarily used for set similarity estimation [2], that have shown promise in fast genomic and metagenomic distance estimation by representing sequences as collections of informative substrings [12]. From locality sensitive hashing literature we know that any sketch over the $\ell^{p}$-norm $(p \in(0,2])$ will automatically lead to sub-quadratic nearest neighbor search by locality sensitive hashing [4]. This provides further motivation for focusing on sketching the edit distance, leaving behind the inherent problems of seed-based methods. To our knowledge, the most accurate sketching method currently available is Ordered MinHash (OMH) [9, which dramatically improves accuracy of $k$-mer MinHash, by considering $k$-mers that are non-adjacent. Therefore, we will compare our approach against the classical $k$-mer MinHash, and Ordered MinHash as the current state-of-the-art.

An especially difficult task for most current methods is the estimation of edit distance for longer, less closely related sequences. To illustrate the main limitations of using $k$-mer statistics in this context, we look at two examples that act as opposing forces in deciding the size of $k$.

Example 1. Consider two sequences $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{A}^{N}$ drawn randomly over the alphabet of size $\# \mathcal{A}=4$, and let $v_{\boldsymbol{x}}, v_{\boldsymbol{y}} \in \mathbb{R}^{4^{k}}$ denote the $k$-mer frequency profile for each sequence. For long sequences with $N \gg 4^{k}, k$-mer frequencies will converge to their mean, that is $4^{-k}$ for all their components, which implies that $\left\|v_{x}-v_{y}\right\| \rightarrow 0$. Therefore, any $k$-mer profile-based method will severely underestimate distance between these two random sequences. In order to avoid this, $k$ has to be restricted to larger values $k \gtrsim \log _{4}(N)$.

Example 2. Now, consider a second pair of sequences, where $\boldsymbol{y}$ is generated by substituting every character of $\boldsymbol{x}$ with $8 \%$ probability. Plugging this value back into our lower bound for $k$, the probability that a $k$-mer would not overlap any substitution equals $(1 / 8)^{\log _{4}(N)}=N^{-1.5}$. Therefore, the total number of $k$-mers not overlapping any mutation will be $N \cdot N^{-1.5}=\left(\frac{1}{N}\right)^{1 / 2}$ which converges to 0 as $N$ grows. Since the sequences share a $92 \%$ similarity in global alignment by design, the frequency vectors are guaranteed to overestimate the distance.

By combining the insights from these examples, we conclude that any value of $k$ will fail with high probability in one of these scenarios for distantly related sequences. Since both results hold for $k$-mer frequency vectors in general, any such statistic on the $k$-mer profile is also bound to fail. This inherent problem with $k$-mers is our main motivation for designing sketches that are more resilient to mutations.

Conceptually, our method can be seen as expectation over all spaced seeds, without any limits on the number of ignore positions. In other words, Tensor Sketching resolves the inherent limitations of any seed-based method by taking an average of all possible seeds, and thereby reducing the risk of selecting a seed that reduces the sensitivity. In statistical terms, this implies that the resulting sketch is a Bayesian admissible estimator.

We will begin by introducing our notation and terminology in Section 2.1 and then proceed to presenting the Tensor Sketch and Tensor Slide sketch methods in Sections 2.2 and 2.3 , respectively. In subsequent Section 3 we first summarize our data generation scheme and then compare our approach to the available state-of-the-art methods. We then discuss applications such as phylogeny reconstruction and graph distance estimation. In the final Section 4 we put our work into context and give an outlook into future work.

## 2 Methods

### 2.1 Preliminaries and Notation

Sets For integer $N,[N]$ denotes the integer set $[N]:=\{1, \ldots, N\}$. For finite set $S$, let $\mathcal{H}_{S}$ denote set of all permutations of indexing $S$

$$
\begin{equation*}
\mathcal{H}_{S}:=\{h: S \mapsto[|S|] \mid\{h(a): a \in S\} \text { is a permutation }\} \tag{1}
\end{equation*}
$$

Vectors Throughout the manuscript bold face fonts are used to distinguish vectors, tensors, and sequences, from scalar variables $\mathbf{a}=\left(a_{1}, a_{2}\right) . \mathbf{1}^{N}$ and $\mathbf{0}^{N}$ are all-ones and all-zeros vectors of length $N$, respectively. We use $i \stackrel{N}{\equiv} j$ when integers $i$ and $j$ are equal modulus $N$. The circular $r$-shift shift $\sin _{r}(\mathbf{a})$, will shift the elements of $\mathbf{a}$ to the $\operatorname{right}_{\operatorname{shift}}^{r}$ ( $\left.\mathbf{a}\right)=\left(a_{N-r+1}, \ldots a_{N}, a_{1}, \ldots, a_{N-r}\right)$, defined formally as

$$
\begin{equation*}
\operatorname{shift}_{r}(\mathbf{a}):=\left(a_{\sigma(i)}\right)_{i \in[N]}, \quad \sigma(i) \stackrel{N}{\equiv} i-r, \sigma:[N] \mapsto[N] \tag{2}
\end{equation*}
$$

where the mapping $\sigma$, circularly shifts indices to the right.

Strings $\mathcal{A}, \mathcal{A}^{N}$, and $\mathcal{A}^{\star}$ denote the alphabet, strings of length $N$, and all strings respectively. $\boldsymbol{x} \circ \boldsymbol{y}$ denotes concatenation, and $|\boldsymbol{x}|$ is the length of $\boldsymbol{x}, \boldsymbol{x}_{i}$ denotes the $i$-th character, $\boldsymbol{x}_{[i: j]}$ is a slice from $i$ th to $j$ th index $\boldsymbol{x}_{[i: j]}:=\boldsymbol{x}_{i} \ldots \boldsymbol{x}_{j}$, referred to as a substring of $\boldsymbol{x}$, while $\boldsymbol{x}_{[\mathrm{i}]}:=\boldsymbol{x}_{i_{1}} \ldots \boldsymbol{x}_{i_{k}}$, is a $k$-ary subsequence (or just subsequence) of $\boldsymbol{x}$, when $\mathbf{i}=\left(i_{1}, \ldots, i_{k}\right)$ is a strictly increasing sequence that $1 \leq i_{1}<\cdots<i_{k} \leq|\boldsymbol{x}|$. For two strings $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{A}^{\star}$, the edit distance $\mathrm{d}_{\mathrm{ed}}(\boldsymbol{x}, \boldsymbol{y})$ denotes the minimum number of edit operations needed to transform one string to the other. It can be defined recursively using $\boldsymbol{x}^{\prime}=\boldsymbol{x}_{[2:|\boldsymbol{x}|]}, \boldsymbol{y}^{\prime}=\boldsymbol{y}_{[2:|\boldsymbol{y}|]}$ as follows

$$
\begin{equation*}
\mathrm{d}_{\mathrm{ed}}(\boldsymbol{x}, \boldsymbol{y}):=\min \left\{\mathrm{d}_{\mathrm{ed}}\left(\boldsymbol{x}^{\prime}, \boldsymbol{y}\right)+1, \mathrm{~d}_{\mathrm{ed}}\left(\boldsymbol{x}, \boldsymbol{y}^{\prime}\right)+1, \mathrm{~d}_{\mathrm{ed}}\left(\boldsymbol{x}^{\prime}, \boldsymbol{y}^{\prime}\right)+\mathbb{1}_{\left\{\boldsymbol{x}_{1} \neq \boldsymbol{y}_{1}\right\}}\right\}, \tag{3}
\end{equation*}
$$

with $\mathrm{d}_{\mathrm{ed}}(\varepsilon, \varepsilon):=0$ as the recursion basis.

Minimum hash sketching For sequence $\boldsymbol{x} \in \mathcal{A}^{N}$, let $X_{i}$ denote the $i$ th $k$-mer $X_{i}:=\boldsymbol{x}_{[i: i+k-1]}$, and $\# X_{i}$ denote its occurrence number $\# X_{i}:=\#\left\{j \leq i: X_{j}=X_{j}\right\}$, and unified $k$-mer refers to the pair $\left(X_{i}, \# X_{i}\right)$.

In MinHash (MH), a random hash is used to sort the $k$-mers, the $k$-mer with the lowest index, with respect to hash functions $h_{1}, \ldots, h_{D} \sim \mathcal{H}_{\mathcal{A}^{k}}$

$$
\begin{equation*}
\boldsymbol{\varphi}_{m h}(\boldsymbol{x}):=\left(X_{j_{1}}, \ldots, X_{j_{D}}\right) \quad j_{r}:=\underset{i}{\operatorname{argmin}}\left\{h_{r}\left(X_{i}\right)\right\} \tag{4}
\end{equation*}
$$

Weighted MinHash (WMH) draws hash functions from unified $k$-mers $h_{1}, \ldots, h_{D} \sim \mathcal{A}^{k} \times[N]$

$$
\begin{equation*}
\boldsymbol{\varphi}_{w m h}(\boldsymbol{x}):=\left(X_{j_{1}}, \ldots, X_{j_{D}}\right) \quad j_{r}:=\underset{i}{\operatorname{argmin}}\left\{h_{r}\left(X_{i}, \# X_{i}\right)\right\} \tag{5}
\end{equation*}
$$

Ordered MinHash (OMH) [9], generalize WMH by sampling from multiple $k$-mers that that appear in the same order in sequence. The sampling uses hashes over the $k$-mers space $h_{1}, \ldots, h_{D} \sim$ $\mathcal{H}_{\mathcal{A}^{k} \times[N]}$, and select indices that are mapped to the $t$ lowest hash values. Formally, for each has function $h_{r}$, we sketch $\left(X_{\sigma_{1}}, \ldots, X_{\sigma_{t}}\right)$, when monotone indices $1 \leq \sigma_{1}<\cdots<\sigma_{t} \leq|X|$ are mapped to the lowest hash values

$$
\begin{equation*}
\left(\boldsymbol{\varphi}_{o m h}(\boldsymbol{x})\right)_{r}:=X_{\sigma_{1}} \circ \cdots \circ X_{\sigma_{t}} \quad\left|\left\{i: h_{r}\left(X_{i}, \# X_{i}\right) \leq h_{r}\left(X_{\sigma_{j}}, \# X_{\sigma_{j}}\right)\right\}\right| \leq t, \text { for all } j \in[t] \tag{6}
\end{equation*}
$$

Finally, we can compute MinHash distance via the Hamming distance $\mathrm{d}_{\mathrm{H}}(X, Y)$, defined as the number of indices that two points $X$ and $Y$ differ $\mathrm{d}_{\mathrm{H}}(X, Y)=\#\left\{i \in[N]: X_{i} \neq Y_{i}\right\}$. Therefore, when $\varphi$ is a MinHash sketch, $\mathrm{d}_{\mathrm{H}}(\varphi(\boldsymbol{x}), \varphi(\boldsymbol{y}))$ becomes an estimator for Jaccard set similarity index $\frac{|X \cap Y|}{|X \cup Y|}$. Therefore, different MinHash methods differ in their choice of set they sketch. Observe that MH, WMH, and OMH are MinHash over $k$-mers, unified $k$-mers, and $t$-tuples of $k$-mers respectively, which in turn implies the equivalence of OMH with WMH for parameter choice of $t=1$.

### 2.2 Tensor Sketch

First, define $\mathcal{I}_{N}^{t}$ to be set of all increasing subsequences of $[N]^{t}$

$$
\begin{equation*}
\mathcal{I}_{N}^{t}:=\left\{\left(\sigma_{1}, \ldots, \sigma_{t}\right) \in[N]^{t}: \text { strictly increasing } \sigma_{1}<\cdots<\sigma_{t}\right\} \tag{7}
\end{equation*}
$$

Without loss of generality let the alphabet be relabeled to $\mathcal{A}:=\{1, \ldots, m\}$, and $m:=\# \mathcal{A}$ be the size. The main idea behind sequence tensor sketch is counting subsequences as opposed to $k$-mers. Define $T_{x}^{(t)}: \mathcal{A}^{t} \rightarrow[0,1]$ to be the probability mass function of a random $t$-ary subsequence of $x$

$$
\begin{equation*}
T_{x}^{(t)}[\boldsymbol{a}]=\mathbb{P}\left(\boldsymbol{x}_{[\mathbf{i}]}=\boldsymbol{a} \mid \mathbf{i} \sim \mathcal{I}_{|\boldsymbol{x}|}^{t}\right) \tag{8}
\end{equation*}
$$

With a slight abuse of notation, we treat $T_{x}^{(t)}$ both as a probability mass function, and an $m^{t}$ dimensional tensor. A simple yet elegant observation links the tensor representation to ordered MinHash. The inner product equals the number of ordered $t$-ary tuples that are shared between $\boldsymbol{x}$ and $\boldsymbol{y}$, normalized by the total number, which appears as the numerator in the Jaccard index. Therefore, OMH and Tensor Sketch are closely related from a theoretical perspective. Then for two arbitrary strings $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{A}^{\star}$, define order $t$ similarity $s^{(t)}(\boldsymbol{x}, \boldsymbol{y})$ and distance $d^{(t)}(\boldsymbol{x}, \boldsymbol{y})$ as

$$
\begin{equation*}
s^{(t)}(\boldsymbol{x}, \boldsymbol{y})=\left\langle T_{\boldsymbol{x}}^{(t)}, T_{\boldsymbol{y}}^{(t)}\right\rangle, \quad \quad d^{(t)}(\boldsymbol{x}, \boldsymbol{y}):=\left\|T_{\boldsymbol{x}}^{(t)}-T_{\boldsymbol{y}}^{(t)}\right\|_{2}^{2} \tag{9}
\end{equation*}
$$

However, naive computation of equations (8) and (9) requires $\binom{N}{t}$ time and $m^{t}$ memory, which becomes prohibitive for even moderate values of these variables. Our tensor sketching scheme, provides a $(1+\epsilon)$-approximation for tensor distance, but only requires $\mathcal{O}\left(\frac{1}{\epsilon^{2}} N t\right)$ time to compute. Furthermore, any downstream analysis will operate on sketches of size $\frac{1}{\epsilon^{2}}$, regardless of the original sequence length.

Since we are interested in preserving the Euclidean norm after sketching, we follow the definition of tensor sketch by Pham and Pagh [14]. For integer $D$, tensor sketch $\Phi: \mathbb{R}^{m^{t}} \rightarrow \mathbb{R}^{D}$ sketches an $m^{t}$-dimensional tensor into $\mathbb{R}^{D}$. Let us assume a series of pairwise independent hash functions $h_{1}, \ldots, h_{t}: \mathcal{A} \rightarrow[D]$, and sign hash functions $s_{1}, \ldots, s_{t}: \mathcal{A} \rightarrow\{-1,1\}$. Moreover, define hash sum $H: \mathcal{A}^{t} \rightarrow[D]$ and hash sign product $S: \mathcal{A}^{t} \rightarrow\{-1,1\}$ as follows

$$
\begin{equation*}
H(\boldsymbol{a}):=\sum_{i}^{t} h_{i}\left(a_{i}\right) \bmod D, \quad S(\boldsymbol{a}):=\prod_{i}^{t} s_{i}\left(a_{i}\right), \quad \boldsymbol{a} \in \mathcal{A}^{t} \tag{10}
\end{equation*}
$$

Finally, tensor sketch $\Phi(T):=\left(\phi_{r}(T)\right)_{r \in[D]}$ is defined as

$$
\begin{equation*}
\phi_{r}(T):=\sum_{\mathbf{a} \in \mathcal{A}^{t}: H(\mathbf{a})=r} S(\mathbf{a}) T[\mathbf{a}] \tag{11}
\end{equation*}
$$

where $T \in \mathbb{R}^{m^{t}}$ is an arbitrary tensor. Crucially, tensor sketch preserves the Euclidean norm in expectation $\mathbb{E}\|\Phi(T)\|_{2}^{2}=\|T\|_{2}^{2}$, and its accuracy increases with sketch size $\operatorname{Var}\|\Phi(T)\|_{2}^{2} \leq \frac{1}{D}\|T\|_{2}^{2}$ which leads to the second moment bound (for proof, see Lemma 7 of Pham and Pagh [14]):

$$
\begin{equation*}
\mathbb{P}\left(\left|\|\Phi(T)\|_{2}^{2}-\|T\|_{2}^{2}\right| \geq \epsilon\|T\|_{2}^{2}\right) \leq \frac{2}{\epsilon^{2} D} \tag{12}
\end{equation*}
$$

Therefore, a sketch size of $D=4 \epsilon^{-2}$ suffices to bound the multiplicative error by $\epsilon$ with probability $\frac{1}{2}$, and sketch size of $D=\mathcal{O}\left(\epsilon^{-2} \log (1 / \delta)\right)$ suffices to boost the probability of success to $1-\delta$, for arbitrary $\delta \in(0,1)$, by taking the median of $\log (1 / \delta)$ independent $\mathcal{O}\left(\epsilon^{-2}\right)$-size sketches.

If we could compute sketches $\Phi\left(T_{x}^{(t)}\right)$ and $\Phi\left(T_{y}^{(t)}\right)$ efficiently, distance computation $\| \Phi\left(T_{x}^{(t)}\right)-$ $\Phi\left(T_{y}^{(t)}\right) \|_{2}^{2}$ merely requires $\mathcal{O}(D)$ time, as opposed to the exponential cost of constructing the ambient space. However, since the tensors $T_{x}^{(t)}$ are never rank-1, the conventional efficient tensor sketching schemes cannot be applied here. Therefore, we have designed an efficient algorithm to compute (11), and provide several extensions of the algorithm in the next section. In summary, for sequence $\boldsymbol{x}$, the tensor sketch $\varphi_{t s}(\boldsymbol{x})$ is computed by applying the tensor sketch on tuple tensor $T_{x}^{(t)}$

$$
\begin{equation*}
\boldsymbol{\varphi}_{t s}(\boldsymbol{x}):=\Phi\left(T_{\boldsymbol{x}}^{(t)}\right) \tag{13}
\end{equation*}
$$

Implicit sketching by dynamic programming First, observe that the $r$ th component of the sketch $\phi_{r}\left(T_{x}^{(t)}\right)$, can be rewritten based on the probability mass of $\left(H\left(\boldsymbol{x}_{[\mathbf{i}]}\right), S\left(\boldsymbol{x}_{[\mathbf{i}]}\right)\right)$, when $\mathbf{i}$ uniformly drawn ordered $t$-tuple $\mathbf{i} \sim \mathcal{I}_{N}^{t}$

$$
\begin{align*}
\phi_{r}\left(T_{x}^{(t)}\right) & =\sum_{b \in\{1,-1\}}(b) \sum_{a \in \mathcal{A}^{t}: H(\boldsymbol{a}) \underline{\equiv}_{=r, S(\boldsymbol{a})=b}} T_{\boldsymbol{x}}^{(t)}[\boldsymbol{a}]  \tag{14}\\
& =\sum_{b \in\{-1,1\}}(b) \mathbb{P}\left(H\left(x_{[\mathbf{i}]}\right)=r, S\left(x_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{|\boldsymbol{x}|}^{t}\right) \tag{15}
\end{align*}
$$

We set out to design a dynamic program that computes probability mass for ordered $t$-tuples, based on probability mass of shorter smaller ordered tuples. Therefore, for $p \in\{1, \ldots, t\}$ let $T_{x}^{(p)}(r, b)$ be the probability mass function of $\left(H_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right), S_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)\right)$, when $\mathbf{i} \sim \mathcal{I}_{p}^{N}$ is a random ordered $p$-tuple over $[N]$, while $H_{p}: \mathcal{A}^{p} \rightarrow[D]$ and $S_{p}: \mathcal{A}^{p} \rightarrow\{-1,1\}$ are partial hash sum and products up to $p$, respectively

$$
\begin{array}{ll}
H_{p}(\boldsymbol{a}):=\sum_{j=1}^{p} h_{j}\left(a_{j}\right) \quad(\bmod D), \quad S_{p}(\boldsymbol{a}):=\prod_{j=1}^{p} s_{j}\left(a_{j}\right), & \boldsymbol{a} \in \mathcal{A}^{p} \\
T_{\boldsymbol{x}}^{(p)}(r, b):=\mathbb{P}\left(H_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=r, S_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{|\boldsymbol{x}|}^{p}\right), & b \in\{-1,1\} \tag{17}
\end{array}
$$

To recursively compute $T_{x_{[1: N]}}^{(p)}(r, b)$ for $p \geq 2$, we separate it into two cases depending on whether the last index is part of the tuple $\left(i_{p}=N\right)$ or $\operatorname{not}\left(i_{p}<N\right)$

$$
\begin{align*}
T_{\boldsymbol{x}_{[1: N]}}^{(p)}(r, b)= & \mathbb{P}\left(H_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=r, S_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{N}^{p}\right)  \tag{18}\\
= & \overbrace{\mathbb{P}_{\mathbf{i}}\left(i_{p}=N\right)}^{z:=} \mathbb{P}\left(H_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=r, S_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{N}^{p}, i_{p}=N\right)+ \\
& \mathbb{P}_{\mathbf{i}}\left(i_{p}<N\right) \mathbb{P}\left(H_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=r, S_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{N}^{p}, i_{p}<N\right)  \tag{19}\\
= & z \mathbb{P}(\sum_{j=1}^{p-1} h_{j}\left(\boldsymbol{x}_{\mathbf{i}_{j}}\right)+\overbrace{h_{p}\left(\boldsymbol{x}_{N}\right)}^{r^{\prime}:=}=r, \prod_{j=1}^{p-1} s_{j}\left(\boldsymbol{x}_{\mathbf{i}_{j}}\right) \overbrace{s_{p}\left(\boldsymbol{x}_{N}\right)}^{b^{\prime}:=}=b \mid \mathbf{i} \sim \mathcal{I}_{N-1}^{p-1})+ \\
& (1-z) \mathbb{P}\left(H_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=r, S_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{N-1}^{p-1}\right)  \tag{20}\\
= & z \mathbb{P}\left(H_{p-1}\left(\boldsymbol{x}_{[\mathbf{i}]}\right) \stackrel{D}{=} r-r^{\prime}, S_{p-1}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b b^{\prime} \mid \mathbf{i} \sim \mathcal{I}_{N-1}^{p-1}\right)+ \\
& (1-z) \mathbb{P}\left(H_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=r, S_{p}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{N-1}^{p}\right)  \tag{21}\\
= & z T_{\boldsymbol{x}_{[1: N-1]}}^{(p-1)}\left(r-r^{\prime} \quad(\bmod D), b b^{\prime}\right)+(1-z) T_{\boldsymbol{x}_{[1: N-1]}}^{p)}(r, b), \tag{22}
\end{align*}
$$

where $z=\frac{p}{N}$, since $\binom{N-1}{p-1} /\binom{N}{p}=\frac{p}{N}$. Defining vector $\mathbf{T}_{x}^{(p)}(b):=\left(T_{x}^{(p)}(r, b)\right)_{r \in[D]}$, the index shift $r-r^{\prime}(\bmod D)$, amounts to a circular $r^{\prime}$-shift of $\mathbf{T}_{\boldsymbol{x}_{[1: N-1]}}^{(p)}(b)$. Hence, the recursion in vector form can be written as
$\diamond$ Recursive insert $(p \in[t])$ :

$$
\left\{\begin{array}{l}
\mathbf{T}_{\boldsymbol{x}_{[1: N]}}^{(p)}( \pm 1)=(1-z) \mathbf{T}_{\boldsymbol{x}_{[1: N-1]}}^{(p)}( \pm 1)+z \operatorname{shift}_{r^{\prime}}\left(\mathbf{T}_{\boldsymbol{x}_{[1: N-1]}}^{(p-1)}\left( \pm b^{\prime}\right)\right)  \tag{23}\\
r^{\prime}:=h_{p}\left(\boldsymbol{x}_{N}\right), \quad b^{\prime}:=s_{p}\left(\boldsymbol{x}_{N}\right), \quad z:=\frac{p}{N}
\end{array}\right.
$$

$\diamond$ Basis $(p=0)$ :

$$
\begin{equation*}
\mathbf{T}_{\boldsymbol{x}}^{(0)}( \pm 1):=(\mathbb{1}\{r \stackrel{D}{\equiv} 0\})_{r \in[D]} \tag{24}
\end{equation*}
$$

Interestingly, the recursion relation for variables in $T_{\boldsymbol{x}_{[1: N]}}^{(p)}(\cdot, \cdot)$ depends only on variables in the previous layer $T_{x_{[1: N-1]}}^{(p)}(\cdot, \cdot)$. Therefore, it is sufficient to store the variables corresponding to layer $i$, which implies a memory complexity of $\mathcal{O}(t D)$. This fact is exploited in Algorithm 1, which is why the subscript $\boldsymbol{x}_{[\mathbf{i}]}$ is drooped from $\mathbf{T}$, and uses $\pm 1$ for brevity ${ }^{4}$. The number of sub-problems needed to compute $\Phi\left(T_{x}^{(t)}\right)$ is exactly $2 N t D$, and the cost of each recursion is $\mathcal{O}(1)$. Therefore, we can compute the sketch with dynamic programming in $\mathcal{O}(\operatorname{DtN})$.

```
Algorithm 1: Sequence Tensor Sketching
    Output : Sequence Tensor Sketch \(\Phi \in \mathbb{R}^{D}\)
    Input \(\quad\) : string \(\boldsymbol{x} \in \mathcal{A}^{N}\),
                            hash functions \(h_{1}, \ldots, h_{t}: \mathcal{A} \rightarrow[D]\),
                sign hashes \(s_{1}, \ldots, s_{t}: \mathcal{A} \rightarrow\{-1,1\}\)
    Parameter: \(D\) : sketch size, \(t\) : subsequence length
    \(\mathbf{T}_{ \pm 1}^{(p)} \leftarrow \mathbf{0}^{D}\) for all \(p \in[t]\)
    \(\mathbf{T}_{ \pm 1}^{(0)} \leftarrow(\mathbb{1}\{r \equiv \underline{\equiv} 0\})_{r \in[D]}\)
    for \(i \leftarrow 1\) to \(N\) do
        for \(p \leftarrow t\) to 1 do
            \(z \leftarrow \frac{p}{i}\)
            \(\mathbf{T}_{ \pm 1}^{(p)} \leftarrow(1-z) \mathbf{T}_{ \pm 1}^{(p)}+z \operatorname{shift}_{h_{p}\left(x_{i}\right)}\left(\mathbf{T}_{ \pm s_{p}\left(x_{i}\right)}^{(p-1)}\right)\)
            end
    end
    \(\Phi \leftarrow \mathbf{T}_{+1}^{(t)}-\mathbf{T}_{-1}^{(t)}\)
```


### 2.3 Tensor Slide Sketch

In this section, we assume that our task involves sketching sequences that have a significant overlap. Instead of applying Algorithm 1 to each sequence separately, we can exploit the fact that many subproblems in the recursion (18) are shared. This idea is exploited in Algorithm2, to compute sketches for all $w$-mers of a reference sequence with length $N$ with $\mathcal{O}\left(N D t^{2}\right)$ time and memory complexity, which is only $t$ times more expensive than Algorithm 1 . The length of the sliding window $w$ does not appear in any complexity term, which could provide significant savings, considering that $t$ is typically much smaller.

The tensor slide sketch, is defined as

$$
\begin{equation*}
\boldsymbol{\varphi}_{t s s}(\boldsymbol{x}):=\left(\boldsymbol{\varphi}_{t s}\left(\boldsymbol{x}_{[i w+1:(i+1) w]}\right)\right)_{i \in\{0,1, \ldots,\lfloor N / w\rfloor\}} \tag{25}
\end{equation*}
$$

where $\varphi_{t s}(\boldsymbol{x})$, is Tensor Sketch $(13)$, the only difference being, that successive sketches are computed dynamically. It is important to note that Tensor Slide Sketch (TSS) only stores $\frac{1}{w}$ of the indices. Because sliding windows that are too overlapping cannot improve our estimate, we can downsample the slide sketch to keep the sketch condense.

## 3 Experimental Results

The primary objective of our tensor sketching framework is to efficiently sketch global alignment scores between sequences. In order to put our method into context, we compare its performance

[^0]```
Algorithm 2: Tensor Slide Sketch
    Output : Sketches \(\Phi_{1}:=\Phi\left(T_{x[1: w]}^{(t)}\right), \ldots, \Phi_{N-w+1}:=\Phi\left(T_{x[N-w+1: N+]}^{(t)}\right)\)
    Input : reference sequence \(\boldsymbol{x} \in \mathcal{A}^{N}\),
                hash functions \(h_{1}, \ldots, h_{t}: \mathcal{A} \rightarrow[D]\),
                sign hashes \(s_{1}, \ldots, s_{t}: \mathcal{A} \rightarrow\{-1,1\}\)
    Parameter: \(D\) : sketch size,
                \(w\) : sliding window,
                \(t\) : tuple length
    \(\mathbf{T}_{ \pm 1}^{(p: q)} \leftarrow \mathbf{0}^{D}\) for all \(1 \leq p \leq q \leq t\)
    \(\mathbf{T}_{ \pm 1}^{(p: q)} \leftarrow(\mathbb{1}\{r \stackrel{D}{\underline{D}} 0\})_{r \in[D]}\) for all \(1 \leq q<p \leq t\)
    for \(i \leftarrow 1\) to \(N\) do
        \(x_{\text {ins }} \leftarrow x_{i}\)
        for all \(1 \leq p \leq q \leq t\) in descending \((q-p)\) do
            \(z \leftarrow \frac{q-p+1}{w}\)
            \(\mathbf{T}_{ \pm 1}^{(p: q)} \leftarrow(1-z) \mathbf{T}_{ \pm 1}^{(p: q)}+z \operatorname{shift}_{h_{q}\left(x_{i n s}\right)}\left(\mathbf{T}_{ \pm s_{q}\left(x_{i n s}\right)}^{(p: q-1)}\right)\)
        end
        if \(i \leq w\) then
        | \(\overline{\text { Continue }}\)
        end
        \(x_{d e l} \leftarrow x_{i-w}\)
        for \(1 \leq p \leq q \leq t\) in ascending \((q-p)\) do
            \(z \leftarrow \frac{q-p+1}{w-q+p}\)
            \(\mathbf{T}_{ \pm 1}^{(p: q)} \leftarrow(1+z) \mathbf{T}_{ \pm 1}^{(p: q)}-z \operatorname{shift}_{h_{q}\left(x_{i n s}\right)}\left(\mathbf{T}_{ \pm s_{q}\left(x_{i n s}\right)}^{(p: q-1)}\right)\)
        end
        \(\Phi_{(i-w+1)} \leftarrow \mathbf{T}_{+1}^{(1: t)}-\mathbf{T}_{-1}^{(1: t)}\)
    end
```

to MH, WMH, and OMH [9], three state-of-the-art sketching techniques (see Section 2 for formal definitions).

### 3.1 Synthetic data generation

As a basis for evaluation, we generated a set of synthetic sequence pairs over the full spectrum of edit distances using Monte Carlo sampling. We draw a reference sequence at random $\boldsymbol{x} \sim \mathcal{A}^{N}$, and for the desired edit distance $d \in[N]$, we generate $\boldsymbol{y}$ by mutating with $d$ edit operations at random $\boldsymbol{x}=: \boldsymbol{x}_{0} \rightarrow \boldsymbol{x}_{1} \rightarrow \cdots \rightarrow x_{d}=: \boldsymbol{y}$, where $\boldsymbol{x}_{i+1}$ is created by selecting a random index from $\left[\left|\boldsymbol{x}_{i}\right|\right]$, and either delete, substitute, or insert a random character at that index. We accept the pair $\boldsymbol{x}, \boldsymbol{y}$, if we have $\mathrm{d}_{\mathrm{ed}}(\boldsymbol{x}, \boldsymbol{y})=d$, and compute the sketches and sketch-based distances for each method. This process is repeated to create the full data set. In total, we generated 1000 sequence pairs of length $N=20000$ over an alphabet of size $\# \mathcal{A}=4$. (See Figure B.1 for a raw presentation of approximate vs. exact distances as scatter plot).

### 3.2 Tensor slide sketch achieves high rank correlation with edit distance

For each sketching method, Spearman's rank correlation with the edit distance, the execution time, and the execution time relative to the edit distance is determined (see Table 1). In particular, Tensor Slide Sketch achieves a correlation of 0.881 while reducing the computation time by $98.6 \%$. The 0 rank correlation for MinHash indicates the fact that very long sequences have 0 distance, as they contain every possible $k$-mer. Furthermore, the AUC are calculated for classification of sequences based on a threshold for their exact edit distance (normalized by length). For all thresholds, TSS clearly outperforms other sketches.

We then assessed the performance and execution time of all considered methods utilizing the synthetically generated data set (Figure 1). We observe that the quality of both MinHash (MH) and Weighted MinHash (WMH) greatly diminishes as the sequences become more distant from

Table 1: 1000 sequence pairs of length $N=20000$ were generated over an alphabet of size $\# \mathcal{A}=4$. with the number of random edit operations, uniformly drawn from $\{0,1, \ldots, N\}$. The time column shows normalized time in microseconds, i.e., total time divided by number of sequences, while the relative time shows the ratio of sketch-based time to the time for computing exact edit distance. As for the the model parameters, embedding dimension is set to $D=50$, and model parameters are (a) MinHash $k=3$, (b) Weighted MinHash $k=3$, (c) Ordered MinHash $k=3, t=3$, (d) Tensor Sketch $t=3$, (e) Tensor Slide Sketch $w=32, t=3$.

|  | Correlation | AUROC $\left(\mathrm{d}_{\mathrm{ed}} \leq \cdot\right)$ |  | Time |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Method | Spearman | 0.10 | 0.20 | 0.5 | Abs. $\left(10^{-6} \mathrm{sec}\right)$ | Rel.(1/ED) |
| MH | 0 | 0.500 | 0.500 | 0.500 | $\mathbf{1 . 3 9 8}$ | 0.001 |
| WMH | 0.361 | 0.795 | 0.746 | 0.670 | 6.497 | 0.004 |
| OMH | 0.466 | 0.917 | 0.817 | 0.700 | 194.195 | 0.120 |
| TS | 0.687 | 0.968 | 0.890 | 0.837 | 45.591 | 0.028 |
| TSS | $\mathbf{0 . 8 8 1}$ | $\mathbf{0 . 9 8 0}$ | $\mathbf{0 . 9 6 5}$ | $\mathbf{0 . 9 3 9}$ | 22.198 | 0.014 |
| ED | - | - | - | - | 1614.040 | 1.000 |

one another (Figure 1a) and the quality of MH and WMH drops with increasing sequence length (Figure 1b). In both plots, the accuracy of OMH and the tensor sketch methods maintains a much higher level. Notably, the increased accuracy of TS and TSS comes only at a very small additional cost in time when compared with MH (Figure 1C), while both methods run much faster than OMH. Lastly, we assessed the relationship of sketching accuracy to the number of embedding dimensions for each method (Figure 1d). Remarkably, TSS achieves a rank correlation of $0.74 \%$ even when using only 10 embedding dimensions.

In summary, TSS produces the most accurate sketches across the board, while introducing a small time and memory overhead when compared with min-hashing. Furthermore, TS and TSS can be computed in a streaming fashion. For example, in order to sketch all 100 -mers of a reference sequence, the time complexity grows only linearly with the length of the reference sequence, while WMH and MH need to be computed separately for each sequence of length $m$, rendering them $m$ times slower for this task.

While sorting creates an iterative bottleneck for all hash-based schemes, the tensor sketch operation is based on a vector shift and sum, which explains the faster execution in practice. Although we have only reported times for single-thread execution, we can expect tensor sketch to achieve a greater boost from multi-threading, as vector shift and sum can be easily vectorized for parallel execution. Moreover, vector operations can be executed in a single step on some CPU or GPU architectures to achieve an even greater boost. Finally, the hash-based methods store $k$-mers, or tuples of $k$-mers as the sketch, which grows with alphabet size and model parameters. In contrast, the tensor sketches can be stored as small integers, namely in a single byte.

### 3.3 Tensor Sketching helps estimate phylogenies

We further explored the task of estimating all pairwise distances of a given set of evolutionary related sequences, resulting in a well structured distance matrix. This task is reminiscent of important bioinformatics applications, such as the binning of metagenomic reads or phylogeny reconstruction. In both cases, the reconstruction of the good distance matrix as a whole, and not only individual components, forms the basis for a good approximation. While there are several ways to formulate this problem, we can simply visualize the exact and approximate matrices and compare them with an overall metric such as Spearman's rank correlation. Figure 2, shows such a distance matrix. The sequences were generated to emulate a phylogeny tree.

### 3.4 Tensor Sketching supports Discretization

While the result of sketching is a rational number, in our applications, these numbers were discretized to take less space. In fact, one can store the sign bit, plus its logarithm with $(1+\epsilon)$ base


Fig. 1: The dataset for these experiments consisted of 10000 sequence pairs independently generated over an alphabet of size $\# \mathcal{A}=4$, and models parameters: MH: $k=3$, WMH: $k=3$, OMH: $k=3, t=3, \mathrm{TS}: t=3, \mathrm{TSS}: w=32, t=3(1 \mathrm{a})$ Area Under the ROC Curve (AUC) distribution, for detection of distance pairs below a threshold sketch-based distances. The x axis, shows which edit distance rank (normalized) is used, and the y axis shows AUC for various sketch-based distances. (1b) For the same dataset as in (1c), the Spearman's rank correlation is plotted against the sequence length (logarithmic scale), 1c) Execution time of each sketching method (y-axis, logscale) as a function of sequence length one (x-axis, log scale). The reported times are normalized, i.e., total sketching time divided by the number of sequences. (1d) Accuracy of various sketching methods as a function of number of embedding dimensions $D$. All methods reach a plateau at some point, after which quality of estimation does improve with sketch size.
rounded to the nearest integer $i$, to ensure an $\epsilon$-bound on the error, as the original number must lie in the interval between $(1+\epsilon)^{i}$ and $(1+\epsilon)^{i+1}$. We used 256 bins, which implies one byte per dimension. In contrast, MH and WMH sketch $k$-mers for each embedding dimension, which will take $k \log (\# \mathcal{A})$ bits per dimension. As OMH sketches $t$ separate $k$-mers, this complexity will even increase to $t k \log (\# \mathcal{A})$ bits.

### 3.5 Tensor Sketch opens up a wide range of applications

Sketching de-Brujin graph distances The slide-sketch algorithm presented in 2, assumes only having access to a stream of characters coming in and that the task involves sketching substrings of the entire sequence. While the task description may sound to be restricted to a linear reference sequence, the same assumptions actually apply to any Euler path on a string graph. For example, if the input is a de-Brujin graph and the task is to sketch all vertices of the graph, any traversal of the graph that walks through adjacent vertices can be mapped to the streaming of incoming and outgoing characters, which will fit the assumptions of the slide sketch algorithm. Consequently, the time complexity of sketching vertices of a de-Brujin graph will grow only as a function of its size, i.e., the number of its vertices, as opposed to the size of each vertex.

Robustness to Global Mutations Tensor Sketch is sensitive to global shifts and transpositions of characters, while $k$-mer statistics are entirely local. To illustrate this, consider case when sequence $\boldsymbol{y}$ is constructed from $\boldsymbol{x}$, by swapping the first and second half of the sequence $\boldsymbol{y}:=\boldsymbol{x}\left[1: \frac{N}{2}\right] \circ \boldsymbol{x}\left[\frac{N}{2}+1: N\right]$. It is evident that $(N-k+1)$ out of $N k$-mers are intact under this block permutation, implying that for sufficiently large $N \gg k$ the $k$-mer profile of $\boldsymbol{y}$ will converge to $\boldsymbol{x}$. This may severely overestimate their alignment score. In contrast, only $2\binom{N / 2}{t} /\binom{N}{t} \simeq 2^{-(t-1)}$ of $t$-ary subsequences remain intact. This greater sensitivity makes tensor sketch an appealing method for applications that such block permutations are likely to occur.


Fig. 2: 1000 sequences of length $n=100$ were generated over an alphabet of size 4 . The phylogeny constructed as a perfectly balanced binary tree, by starting from a single sequence, and at each generation mutate every sequence in the pool and add it to the pool. Children were generated with a point-mutation rate of 0.1. Subplot (a) shows the pairwise edit distances normalized by length $n$, while subplots (b)-(f) are approximate sketch-based distances normalized by embedding dimension $m=128$

## 4 Discussion

We presented tensor sketching, a method for estimating sequence similarity for biological sequences. We demonstrated that Tensor Slide Sketch achieves a high Spearman's rank correlation, but runs an order of magnitude faster than computing the exact alignment. When compared with other sketching methods, Tensor Sketch and Tensor Slide Sketch both achieve a more accurate preservation of the order between distances than MinHash, Weighted MinHash, and Ordered MinHash.

Hyper-parameters The results presented in Section 3 have focused on a fixed set of parameters for the models. However, our results are fairly robust to the choice of parameters. Namely, varying the size of $k$-mers will be beneficial to hash-based sketches for small edit distances, while it will hurt the performance for more distant pairs. Hence, Spearman's rank correlation will not substantially improve. While tensor sketch achieves the best accuracy in our set of experiments, a comprehensive study of different parameter regimes and settings under which they will be beneficial merits further investigation. In particular, min-hash methods typically show good performance at detecting small distances, as it is evident in our AUROC plot in Figure 1. This suggests the possibility of having a hybrid-version of tensor sketching and ordered MinHash to combine the best features of the two. Moreover, such a study could shed more light on the connections between min-hashing and tensor sketching.

Bayesian estimation with tensor sketching Spaced $k$-mers were introduced motivated by applications involving distantly related sequences. If we are to allow for $i$ ignore positions, there will be $\binom{k}{i} \approx(k / i)^{i}$ possible patterns. Considering all such patterns is clearly not an option if $i$ is anything larger but a small constant. Not only is it highly non-trivial to search for the optimal pattern in this large combinatorial space, but any incremental step for pattern optimization also has to be repeated if different tasks require a different set of patterns. Seen from this angle,

Tensor Sketch can be interpreted as an average of all possible spaced $t$-mers patterns with $(N-t)$ ignore positions, while the sketching step avoids the exponential cost of explicitly representing all combinatorial options.

The core idea of tensor sketching, that is to average over all possible seed patterns, can be alternatively viewed as Bayesian estimation. This implies that our estimate is admissible, i.e., no other estimator can outperform it across the board. This provides some intuition into why our approach achieves an acceptable or better accuracy over the entire range of edit distances. In other words, the risk of selecting a bad seed or size of $k$-mer is minimized by taking the mean over all possible seeds. While this corresponds to a uniform prior on our seeds, one can introduce a non-uniform prior by setting different weights for vertices on our tensors, namely to penalize more number of gaps in the pattern. Therefore, we can strengthen or weaken the contribution of individual patterns. Weighted average can help to reduce background statistics into the distance metric, as well as introduce nonuniform mismatch penalties. These variations and extensions come naturally within the tensor sketching framework, in contrast with hash-based approaches that require heuristic steps.

Conclusion This work's main contributions are the introduction of two tensor-based sketching methods to sequence similarity estimation: Tensor Sketch, providing an efficient algorithm to compute it, and Tensor Slide Sketch, a streaming version of the algorithm. Our results indicate that the tensor sketching method and its applications open up exciting research directions to explore for the bioinformatics community. The main advantages of tensor sketching compared with hashbased methods are that it 1) can run in a streaming fashion, 2) achieves much higher accuracy, and 3 ) runs much faster in practice.

## Acknowledgements

We would like to thank Kjong-Van Lehmann and Ximena Bonilla for their constructive feedback on the manuscript. AJ was funded by ETH Zurich core funding to GR. AK was partially funded by PHRT Project \#106 and the Swiss National Science Foundation Grant No. 407540_167331 "Scalable Genome Graph Data Structures for Metagenomics and Genome Annotation" as part of Swiss National Research Programme (NRP) 75 "Big Data".

## References

1. Apostolico, A., Denas, O.: Efficient tools for comparative substring analysis. Journal of Biotechnology 149(3), 120-126 (2010). DOI 10.1016/J.JBIOTEC.2010.05.006. URL https://www.sciencedirect.com/science/ article/pii/S0168165610002403?via\%3Dihub
2. Broder, A.Z.: On the resemblance and containment of documents. Proceedings of the International Conference on Compression and Complexity of Sequences pp. 21-29 (1997)
3. Hahn, L., Leimeister, C.A., Ounit, R., Lonardi, S., Morgenstern, B.: rasbhari: Optimizing Spaced Seeds for Database Searching, Read Mapping and Alignment-Free Sequence Comparison. PLoS computational biology 12(10), e1005,107 (2016). DOI 10.1371/journal.pcbi.1005107. URL http://www.ncbi.nlm.nih.gov/pubmed/ 27760124 http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5070788
4. Kulis, B., Grauman, K.: Kernelized locality-sensitive hashing for scalable image search. In: 2009 IEEE 12th international conference on computer vision, pp. 2130-2137. IEEE (2009)
5. Leimeister, C.A., Boden, M., Horwege, S., Lindner, S., Morgenstern, B.: Fast alignment-free sequence comparison using spaced-word frequencies. Bioinformatics (Oxford, England) 30(14), 1991-9 (2014). DOI 10.1093/ bioinformatics/btu177. URL http://www.ncbi.nlm.nih.gov/pubmed/24700317http://www.pubmedcentral. nih.gov/articlerender.fcgi?artid=PMC4080745
6. Leimeister, C.A., Morgenstern, B.: Kmacs: the k-mismatch average common substring approach to alignmentfree sequence comparison. Bioinformatics (Oxford, England) 30(14), 2000-8 (2014). DOI 10.1093/ bioinformatics/btu331. URL http://www.ncbi.nlm.nih.gov/pubmed/24828656http://www.pubmedcentral. nih.gov/articlerender.fcgi?artid=PMC4080746
7. Ma, B., Tromp, J., Li, M.: PatternHunter: faster and more sensitive homology search. Bioinformatics 18(3), 440-445 (2002). DOI 10.1093/bioinformatics/18.3.440. URL http://www.ncbi.nlm.nih.gov/pubmed/ 11934743https://academic.oup.com/bioinformatics/article-lookup/doi/10.1093/bioinformatics/18. 3.440
8. Mak, D.Y., Benson, G.: All hits all the time: parameter-free calculation of spaced seed sensitivity. Bioinformatics 25(3), 302-308 (2009). DOI 10.1093/bioinformatics/btn643. URL https://academic.oup.com/ bioinformatics/article-lookup/doi/10.1093/bioinformatics/btn643
9. Marcais, G., Deblasio, D., Pandey, P., Kingsford, C.: Locality-sensitive hashing for the edit distance. Bioinformatics 35(14), i127-i135 (2019). DOI 10.1093/bioinformatics/btz354
10. Noé, L.: Best hits of 11110110111: model-free selection and parameter-free sensitivity calculation of spaced seeds. Algorithms for molecular biology : AMB 12, 1 (2017). DOI 10.1186/s13015-017-0092-1. URL http://www.ncbi. nlm.nih.gov/pubmed/28289437http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5310094
11. Noe, L., Martin, D.E.: A coverage criterion for spaced seeds and its applications to support vector machine string kernels and k-mer distances. Journal of Computational Biology 21(12), 947-963 (2014). DOI 10.1089/ cmb.2014.0173. URL http://www.ncbi.nlm.nih.gov/pubmed/25393923http://www.pubmedcentral.nih.gov/ articlerender.fcgi?artid=PMC4253314
12. Ondov, B.D., Treangen, T.J., Melsted, P., Mallonee, A.B., Bergman, N.H., Koren, S., Phillippy, A.M.: Mash: Fast genome and metagenome distance estimation using MinHash. Genome Biology 17(1), 1-14 (2016). DOI 10.1186/s13059-016-0997-x. URL http://dx.doi.org/10.1186/s13059-016-0997-x
13. Onodera, T., Shibuya, T.: The Gapped Spectrum Kernel for Support Vector Machines. In: International Workshop on Machine Learning and Data Mining in Pattern Recognition, pp. 1-15. Springer, Berlin, Heidelberg (2013). DOI 10.1007/978-3-642-39712-7\{\_\}1. URL http://link.springer.com/10.1007/ 978-3-642-39712-7_1
14. Pham, N., Pagh, R.: Fast and scalable polynomial kernels via explicit feature maps. In: Proceedings of the 19 th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 239-247 (2013)
15. Röhling, S., Dencker, T., Morgenstern, B.: The number of k-mer matches between two DNA sequences as a function of k. bioRxiv p. 527515 (2019). DOI 10.1101/527515. URL https://www.biorxiv.org/content/10. $1101 / 527515 \mathrm{v} 2$
16. Sims, G.E., Jun, S.R., Wu, G.A., Kim, S.H.: Whole-genome phylogeny of mammals: Evolutionary information in genic and nongenic regions. Proceedings of the National Academy of Sciences of the United States of America 106(40), 17,077-17,082 (2009). DOI 10.1073/pnas.0909377106. URL http://www.ncbi.nlm.nih.gov/pubmed/ 19805074http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2761373
17. Sims, G.E., Kim, S.H.: Whole-genome phylogeny of Escherichia coli/Shigella group by feature frequency profiles (FFPs). Proceedings of the National Academy of Sciences of the United States of America 108(20), 8329-34 (2011). DOI 10.1073/pnas.1105168108. URL http://www.ncbi.nlm.nih.gov/pubmed/21536867http://www. pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3100984
18. Ulitsky, I., Burstein, D., Tuller, T., Chor, B.: The Average Common Substring Approach to Phylogenomic Reconstruction. Journal of Computational Biology 13(2), 336-350 (2006). DOI 10.1089/cmb.2006.13.336. URL http://www.liebertpub.com/doi/10.1089/cmb.2006.13.336

## Appendix A Tensor Slide Sketch Algorithm

The main idea is to extend the data structures, to dynamically append the string or remove a character from the beginning. The procedure is detailed in Algorithm 2 (See Section A for derivation of recursive step in detail). In order to implement an iterative algorithm, one can roll back the recursive calls and execute them in the reverse order. Namely for deletion, recursive relation for $\mathbf{T}_{\boldsymbol{x}_{[2: w+1]}}^{(p: q)}(\cdot)$ depends on $\mathbf{T}_{\boldsymbol{x}_{[1: w+1]}}^{(p: q)}(\cdot)$, and $\mathbf{T}_{\boldsymbol{x}_{[2: w]}}^{(p+1: q)}(\cdot)$, implying that updates for $\mathbf{T}_{\boldsymbol{x}_{[2: w+1]}}^{(p+1: q)}(\cdot)$ must precede updates for $\mathbf{T}_{\boldsymbol{x}_{[2: w+1]}}^{(p: q)}(\cdot)$, imposing a partial order corresponding to tuple length $(q-p+1)$. This in turn corresponds with the random subsequence length. For insertion recursion (18), subproblems correspond to shorter subsequence lengths, while for deletion recursion (31), sub-problems correspond to longer subsequence lengths, which impose an ascending and descending dependency on subsequence length, respectively.

For all $1 \leq p \leq q \leq t$, let $H_{p: q}: \mathcal{A}^{q-p+1} \rightarrow[D]$ and $S_{p: q}: \mathcal{A}^{q-p+1} \rightarrow\{-1,1\}$ be partial hash sum and products from $p$ up to $q$, respectively, and define $T_{x}^{(p: q)}(r, b)$ be the probability mass function of $\left(H_{p: q}\left(\boldsymbol{x}_{[\mathbf{i}]}\right), S_{p: q}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)\right)$, when $\mathbf{i}$ is a uniformly drawn ordered $(q-p+1)$-tuple.

$$
\begin{array}{ll}
H_{p: q}(\boldsymbol{a}):=\sum_{j=p}^{q} h_{j}\left(a_{j}\right) \quad(\bmod D), \quad S_{p: q}(\boldsymbol{a}):=\prod_{j=p}^{q} s_{j}\left(a_{j}\right), & \boldsymbol{a} \in \mathcal{A}^{q-p+1} \\
T_{x}^{(p: q)}(r, b):=\mathbb{P}\left(H_{p: q}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=r, S_{p: q}\left(\boldsymbol{x}_{[\mathbf{i}]}\right)=b \mid \mathbf{i} \sim \mathcal{I}_{|x|}^{q-p+1}\right), & b \in\{-1,1\} \tag{27}
\end{array}
$$

The recursive insertion is analogous to (18), as we can derive $T_{x_{[1: w]}^{(p: q)}}^{(r, b) \text { based on the previous }}$ layer $T_{\boldsymbol{x}_{[1: w-1]}}^{(p: q)}(\cdot, \cdot)$ and smaller problems $T_{\boldsymbol{x}_{[1: w-1]}}^{(p: q-1)}(\cdot, \cdot)$. Recursive deletion is equivalent to rolling back insertion of the first character. We can derive $T_{x_{[1: w+1]}}^{(p: q)}(r, b)$ based on $T_{x[2: w+1]}^{(p+1: q)}\left(r^{\prime}, b^{\prime}\right)$ and $T_{x_{[2: w+1]}}^{(p: q)}(r, b)$, by separating it into two cases when the random tuple starts at the first index being in the tuple

$$
\begin{equation*}
T_{\boldsymbol{x}_{[1: w+1]}}^{(p: q)}(r, b)=(1-t) T_{\boldsymbol{x}_{[2: w]}}^{(p: q)}(r, b)+t T_{\boldsymbol{x}_{[2: w]}}^{(p+1: q)}\left(r-r^{\prime}(\bmod D), b^{\prime}\right), \quad t:=\mathbb{P}\left(\mathbf{i}_{1}=1 \mid \mathbf{i} \sim \mathcal{I}_{w+1}^{q-p+1}\right) \tag{28}
\end{equation*}
$$

where $t=\frac{p-q+1}{w+1}$, analogous to the insertion only case, and $r^{\prime}:=r-h_{p}\left(\boldsymbol{x}_{1}\right)$ and $b^{\prime}:=b s_{p}\left(\boldsymbol{x}_{1}\right)$ are recursive hashes as before. We can rearrange this identity to compute $T_{x_{[2: w]}^{(p: q)}}^{(r, b) \text { based on other }}$ terms

$$
\begin{equation*}
T_{x_{[2: w+1]}}^{(p: q)}(r, b)=\frac{1}{1-t} T_{\boldsymbol{x}_{[1: w+1]}}^{(p: q)}(r, b)-\frac{t}{1-t} T_{\boldsymbol{x}_{[2: w+1]}}^{(p+1: q)}\left(r-r^{\prime}(\bmod D), b b^{\prime}\right) \tag{29}
\end{equation*}
$$

Defining vector $\mathbf{T}_{\boldsymbol{x}}^{(p: q)}(b):=\left(T_{\boldsymbol{x}}^{(p: q)}(r, b)\right)_{r \in[D]}$, and calculating $\frac{t}{1-t}=\frac{q-p+1}{w-q+p}$, we can write the recursion in the vector form more concisely as:
$\diamond$ Recursive delete $(p \leq q)$ :

$$
\left\{\begin{array}{l}
\mathbf{T}_{\boldsymbol{x}_{[2: w+1]}}^{(p: q)}( \pm 1)=(1+z) \mathbf{T}_{\boldsymbol{x}_{[1: w]}}^{(p: q)}( \pm 1)-z \operatorname{shift}_{r^{\prime}}\left(\mathbf{T}_{\boldsymbol{x}_{[2: w]}}^{(p+1: q)}\left( \pm b^{\prime}\right)\right)  \tag{30}\\
b^{\prime}:=s_{p}\left(\boldsymbol{x}_{1}\right), \quad r^{\prime}:=h_{p}\left(\boldsymbol{x}_{1}\right), \quad z:=\frac{q-p+1}{w}
\end{array}\right.
$$

$\diamond$ Recursive insert $(p \leq q)$ :

$$
\left\{\begin{array}{l}
\mathbf{T}_{\boldsymbol{x}_{[1: w]}}^{(p: q)}( \pm 1)=(1-z) \mathbf{T}_{\boldsymbol{x}_{[1: w-1]}}^{(p: q)}( \pm 1)+z \operatorname{shift}_{r^{\prime}}\left(\mathbf{T}_{\boldsymbol{x}_{[1: w]}}^{(p: q-1)}\left( \pm b^{\prime}\right)\right)  \tag{31}\\
b^{\prime}:=s_{q}\left(x_{N}\right), \quad r^{\prime}:=h_{q}\left(x_{N}\right), \quad z:=\frac{q-p+1}{w-q+p}
\end{array}\right.
$$

$\diamond$ Basis $(q<p)$ :

$$
\begin{equation*}
\mathbf{T}_{x}^{(p: q)}(b):=(\mathbb{1}\{r \stackrel{D}{\equiv} 0\})_{r \in[D]} \tag{32}
\end{equation*}
$$



Fig. B.1: Normalized sketch distance ranks versus normalized edit distance ranks. 50000 sequence pairs of length $N=256$ were generated over $\# \mathcal{A}=4$ alphabets. One sequence was generated randomly, and the second was mutated randomly, to generate a spectrum of edit distances. Subplot (a-e) show rank of sketch-based distances vs. rank of edit distances, both normalized by the total number of pairs. The embedding dimension for all models, and parameters for each model are $D=100$, and parameters models are (a) MinHash $k=3$, (b) Weighted MinHash $k=3$, (c) Ordered MinHash $k=3, t=3$, (d) Tensor Sketch $t=3$, (e) Tensor Slide Sketch $w=32, t=3$.

## Appendix B More results

Figure B.1 shows individual pair of sequences as a scatter plot, while the solid lines indicate median and quartiles. As the Spearman's rank correlation indicates, that rank of edit distances are better preserved by the two tensor sketching methods, and more so by tensor slide sketch. This difference is due to the fact that hash-based sketches are sensitive to small edit distances, namely below 0.1 (normalized to length), they plateau more quickly.

We can evaluate these methods more quantitatively, by requiring them to solve a decision edit distance problem. In the decision version of the problem, the task is to decide which points are bellow a certain threshold and which ones are not. Therefore, the performance of each method can be depicted by ROC curve, which are plotted for four different thresholds in Figure B.2, with the dataset remains the same as Figure B.1. The main takeaway of ROC plots is that while hashbased methods are effective at very low thresholds, they lose their sensitivity for higher thresholds, and their sensitivity declines as task moves away from comparing very similar sequence pairs. Interestingly, tensor slide sketch seems to perform nearly as optimal in the lower threshold, but outperforms other methods on most lower thresholds significantly.


Fig. B.2: 10000 sequence pairs of length $n=100$ were generated over an alphabet of size 4 . with the number of random edit operations, uniformy drawn from $\{0,1\}$


[^0]:    ${ }^{4}$ The plus-minus signs are shorthand for two equations: one by taking all to be + , and another taking all to be - , but not any other combination, for example $\sin (a \pm b)=\sin (a) \cos (b) \pm \sin (b) \cos (a)$. If there are 干 in the equation, their order for execution must the reversed, for example $\cos (a \pm b)=\cos (a) \cos (b) \mp \sin (a) \sin (b)$.

