How to optimally sample a sequence for rapid analysis

Martin C. Frith[1,2,3] Jim Shaw[4] John L. Spouge[5]

[1]Artificial Intelligence Research Center, AIST, Japan. [2]Graduate School of Frontier Sciences, University of Tokyo, Japan. [3]Computational Bio Big-Data Open Innovation Laboratory, AIST, Japan. [4]Department of Mathematics, University of Toronto, Canada. [5]National Center for Biotechnology Information, National Library of Medicine, National Institutes of Health, USA.

Abstract

We face an increasing flood of genetic sequence data, from diverse sources, requiring rapid computational analysis. Rapid analysis can be achieved by sampling a subset of positions in each sequence. Previous sequence-sampling methods, such as minimizers, syncmers, and minimally-overlapping words, were developed by heuristic intuition, and are not optimal.

We present a sequence-sampling approach that provably optimizes sensitivity for a whole class of sequence comparison methods, for randomly-evolving sequences. It it likely near-optimal for a wide range of alignment-based and alignment-free analyses. For real biological DNA, it increases specificity by avoiding simple repeats. Our approach generalizes universal hitting sets (which guarantee to sample a sequence at least once), and polar sets (which guarantee to sample a sequence at most once). This helps us understand how to do rapid sequence analysis as accurately as possible.

Introduction

Let us start with a curious puzzle. Suppose we randomly pick one of the 4^5 DNA sequences of length 5. Is the probability that it contains the substring `acgt` equal to the probability that it contains `aaaa`?

There seems to be a strong intuition that the answer is "yes": let us check. There are four length-5 sequences that start with `acgt`, and another four that end with `acgt`, for a total of eight. On the other hand, there are seven sequences containing `aaaa`: `aaaaa` plus three others starting with `aaaa` and three others ending with `aaaa`. It is remarkable that such a simple puzzle can be so counterintuitive.

This curious property of sequences is not just a shallow trick, because it can be leveraged into useful methods. For example, it is the basis of spaced seeds [15]. Here, we build upon it a theory of optimal sequence-sampling, for rapid analysis of big data.

The main way of analyzing genetic sequences is by comparing them to each other. For large data, this is usually done via "seeds", by which we mean simple similarities that can be found quickly. The simplest seeds are fixed-length exact matches, but they can also be inexact [1, 15, 20] and/or variable-length [7]. These seeds are used in diverse ways to infer relationships between sequences.

The seed-and-extend approach extends seeds into detailed alignments. The simplest "one hit" method attempts extension from each seed [1]. A "two hit" alternative triggers extension from a pair of nearby, non-overlapping seeds [2]. Another "coverage" method triggers extension from a group of possibly-overlapping seeds, if the number of letters covered by the seeds exceeds a threshold [20, 3, 21, 19].

There are also alignment-free approaches [27], which often use seeds, also called "micro-alignments" [29]. These methods are used to infer evolutionary relationships, e.g. by estimating evolutionary distances between sequences [5, 17]. For example, they are used to classify sequences taxonomically, based on seed matches to a database of reference genomes [23, 28], which is phylogenetic placement [13].

For faster or larger-scale sequence comparison, per-

haps the most promising approach is sampling, for example just use seeds starting with `a`. Sampling methods include: anchor strings [16], minimizers [25, 24], syncmers [9], and minimally-overlapping words [11]. Those studies showed that some sampling methods are better than others in some situations.

Previous sampling methods, however, were developed by heuristic intuition, and are unlikely to be optimal. It is also possible that a sampling method is superior in some scenarios but not others, e.g. for closely- but not distantly distantly-related sequences, or for alignment-based but not alignment-free methods. Thus, we lack a non-heuristic criterion to optimize sampling.

Here we develop sampling methods, for exact-match seeds, which provably optimize sensitivity for one-hit-based sequence comparison, for randomly-evolving sequences. We argue that they are likely near-optimal in a wide range of alignment-based and alignment-free scenarios. We then show evidence that they work well for real biological DNA.

Illustration of our sampling approach

The main idea is that related sequences have runs of overlapping matches (fig. 1), and we maximize the probability of sampling these runs. To make things simple, we assume a random DNA sequence of independent, equally-probable bases, evolving by random substitutions. Let us consider two sampling schemes: matches starting with `ry`, and matches starting with `rr`. (`r` means puRine: `a` or `g`, and `y` means pYrimidine: `c` or `t`.) Each scheme samples 1 in 4 matches on average. A run of, say, 2 overlapping matches is hit when `ry` (or `rr`) occurs at the start of either match, i.e. when it occurs in a sequence of length 3. A quarter of all length-3 sequences start with `ry`, and a *different* quarter end with `ry` (because it's impossible to start *and* end with `ry`), so the total hitting probability is $\frac{1}{4} + \frac{1}{4} = \frac{1}{2}$. In contrast, $\frac{1}{4}$ of length-3 sequences start with `rr` and an *overlapping* quarter end with `rr`, so the total hitting probability is lower ($\frac{3}{8}$). Thus, `ry` is more likely than `rr` to sample a run of 2 overlapping matches.

In fact, `ry` is more likely than `rr` to sample a run of x overlapping matches, for any $x \geq 2$. This means

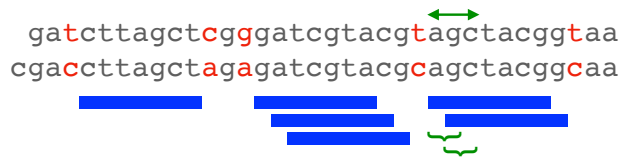

Figure 1: Example of related sequences, with length-7 exact matches (blue bars). The matches occur in runs of size 1, 3, and 2. If we sample matches starting with `ry`, we "hit" the size-2 run if `ry` is at either of the sites marked by braces, in other words, if `ry` occurs in the length-3 sequence marked by arrows.

that `ry` is more likely than `rr` to sample at least one match between two related sequences. Thus, `ry` is more sensitive, for one-hit-based sequence comparison. Moreover, it is more sensitive for any level of sequence divergence: for closely- or distantly-related sequences.

Methods

Word-based sampling

A word-based sampling scheme specifies a set Q of length-k words, e.g. $Q = \{\texttt{rrry}, \texttt{ryrr}, \texttt{ryyr}, \texttt{yyyr}\}$, and selects the positions in a sequence where those words occur. This can be used to select seeds starting at those positions. We will focus on exact-match seeds of fixed length $m \geq k$. Some previous publications (e.g. [9]) have focused on the case $m = k$, so that both words and matches are k-mers, but here words are k-mers and matches are m-mers. Fundamentally, we are sampling positions in a sequence, and what we do with those positions is secondary.

Word-based sampling can be implemented with a lookup table of size $|A|^k$, where $|A|$ is the alphabet size, to look up whether any word is in the set. This is fast and practical, as long as the table is small enough to fit in the computer's fast cache memory.

We use words made of `r` and `y`, for these reasons:

- In biological DNA, `r` and `y` have 50:50 frequencies, whereas `a`, `c`, `g`, and `t` have varying frequencies. This makes the frequency of `r`/`y` words more predictable.
- Vertebrates have a high rate of `cg` → `tg`/`ca` sub-

stitutions [31]. This `cg` effect does not fit our model of independent letters with uniform substitution probability, but it does fit the model when the sites are encoded as `r/y` [4].

- A reduced alphabet makes it easier to use longer words, because the lookup table is smaller. Longer words can help to optimize sampling, by constraining the spacing between sampled positions.
- `r/y` word-based sampling is compatible with inexact seeds that allow `a`↔`g` and `c`↔`t` substitutions at some positions: such seeds often work well because these substitutions are frequent [20].

Run-hitting probabilities

We wish to know the probability that a word-based sampling scheme hits a run of x consecutive start positions (fig. 1). This equals the probability that any word in our set occurs in a sequence of length $x+k-1$. This can be calculated by dynamic programming. If w is any length-k string, define $C(i, w)$ to be the number of possible length-i sequences that end with w and have no substring that is in our set Q. The base case is:

$$C(k, w) = \begin{cases} 0 & \text{if } w \in Q \\ 1 & \text{if } w \notin Q \end{cases} \tag{1}$$

If we define w' to be w with its final letter removed, $\oplus$ is string concatenation, and A is the alphabet:

$$C(i+1,\, w) = \begin{cases} 0 & \text{if } w \in Q \\ \sum_{a \in A} C(i,\, a \oplus w') & \text{if } w \notin Q \end{cases} \tag{2}$$

Finally, from the definition of $C(i, w)$, the probability of any word in our set occurring in x overlapping positions is:

$$H_x = 1 - \frac{\sum_w C(x+k-1, w)}{|A|^{x+k-1}} \tag{3}$$

Upper bound

The run-hitting probabilities have a simple upper bound [26]. This is illustrated by the `ry` example (fig. 1): the scheme samples 1 in s positions ($s = 4$ in this case), the probability of hitting a size-2 run is at most $2 \times (1/s)$, and this upper bound is achieved when it's impossible to sample both positions. In general, the probability of hitting a size-x run is at most: $\min(x/s,\, 1)$.

Optimizing run-hitting probabilities

We wish to find a set of words with high run-hitting probabilities. We will try to maximize the *average* run-hitting probability, for runs of size $1, 2, 3, \ldots u$. When u is very large, this is equivalent to minimizing the variance of distance between consecutive sampled positions (see the Supplement).

We used two optimization methods: exact and heuristic. In both cases, we fix the word length k and the number of words n. The exact method evaluates all possible sets of n length-k words, and outputs all sets that have maximum average run-hitting probability. This is feasible only when k and n are small; we used it when $n \leq 8$. The heuristic method uses simulated annealing (see the Supplement): it does not guarantee to find the maximum average run-hitting probability. These methods are available in `maxhit-words` from `https://gitlab.com/mcfrith/noverlap`.

Reversing to minimize `yr` count

So far, we have assumed that differences between related sequences are random and sequence-independent. This is not really true: in particular, vertebrates have a high rate of `cg` → `tg`/`ca` substitutions [31]. This suggests that words with fewer `yr`s may be better. As a first step in this direction, we reversed all words in our set if doing so would reduce the total `yr` count. Such reversal does not change the run-hitting probabilities calculated as above.

Syncmer run-hitting probabilities

These were calculated using `distance-distribution-parametrized-syncmer.py` from `https://tinyurl.com/frith-ISCB-V-Taiwan`, which builds on "parametrized syncmer" concepts [8].

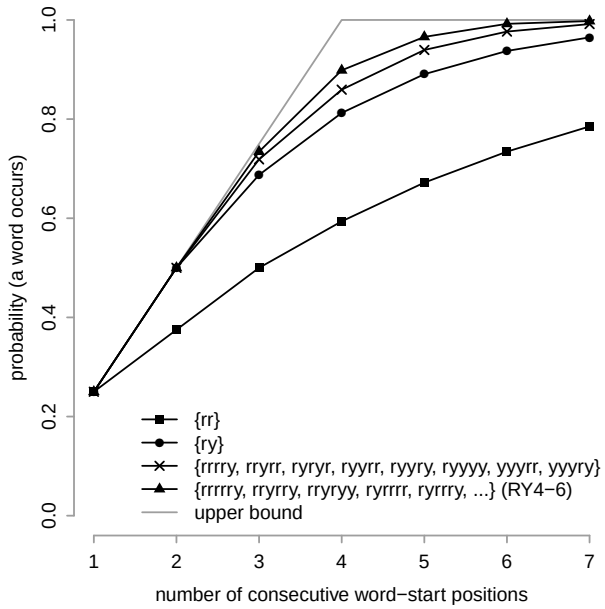

Figure 2: For each of four word sets, probability of occurring at least once in a sequence of random, independent, equally-probable `r` and `y`.

Results

Run-hitting probabilities

To start with our introductory example of `ry` versus `rr`, we calculated the probabilities of `ry` or `rr` occurring in runs of overlapping positions (fig. 1). `ry` has higher probability of occurring in a run of between 2 and at least 7 positions (fig. 2). This indicates that matches starting with `ry` between related sequences are more likely to occur than matches starting with `rr`.

We then calculated run-hitting probabilities for some word sets found previously by heuristic criteria [11]. A set of eight length-5 words (fig. 2), which also samples 1 in 4 positions, has higher run-hitting probabilities than `ry`. Another set of sixteen length-6 words, named `RY4-6`, has even higher run-hitting probabilities (fig. 2). The eight 5-mers have the maximum average run-hitting probability, for run size 1 to u=7, among all sets of eight 5-mers (according to our exact optimization method). This shows that longer words enable higher run-hitting probabilities.

Every s-th sampling is optimal

A simple sampling method is to use matches whose start coordinate, in one of the two sequences, is divisible by s, for example $s = 4$. This scheme hits a run of $x < s$ overlapping matches with probability x/s, and $x \geq s$ matches with probability 1. Thus, it perfectly achieves the best-possible upper bound run-hitting probabilities.

This scheme has a disadvantage: it checks for matches starting every s-th position in just one sequence, but starting at every possible position in the other sequence. In constrast, if we use e.g. matches starting with `a`, we can rapidly skip all non-`a` positions in both sequences.

Another idea is to use matches whose start coordinates, in both sequences, are divisible by $\sqrt{s}$, which samples 1 in s matches on average. This scheme hits runs of size $x \leq \sqrt{s}$ with optimal probability x/s. But for longer runs the hitting probability flatlines at $1/\sqrt{s}$.

Optimizing run-hitting probabilities

We sought word sets whose run-hitting probabilities are as high as possible, for sparsity $s = 4$, 8, 16, and 32. If we use `r`/`y` words of length k, the number of words must be $n = 2^k/s$. Longer words can achieve higher hitting probabilities, but are harder to optimize because there are more possible word sets. So we compromized by using heuristic optimization with $n = 128$ (table 1). We sought word sets with maximum average hitting probability, for run sizes from 1 to $u = 2s - 1$. This upper value u is arbitrary, but covers the range where the probabilities are most different from the upper bound (fig. 2, 3).[1]

We found word sets with higher run-hitting probabilities, for each sparsity, than the best word sets we found previously (fig. 3). The names indicate sparsity and word length, e.g. `RY4-9` has sparsity 4 and word length 9. The improvement is small but clear.

[1] Actually, we set a hard upper limit on u: $u + k - 1 \leq 64$, because our software does not handle higher values.

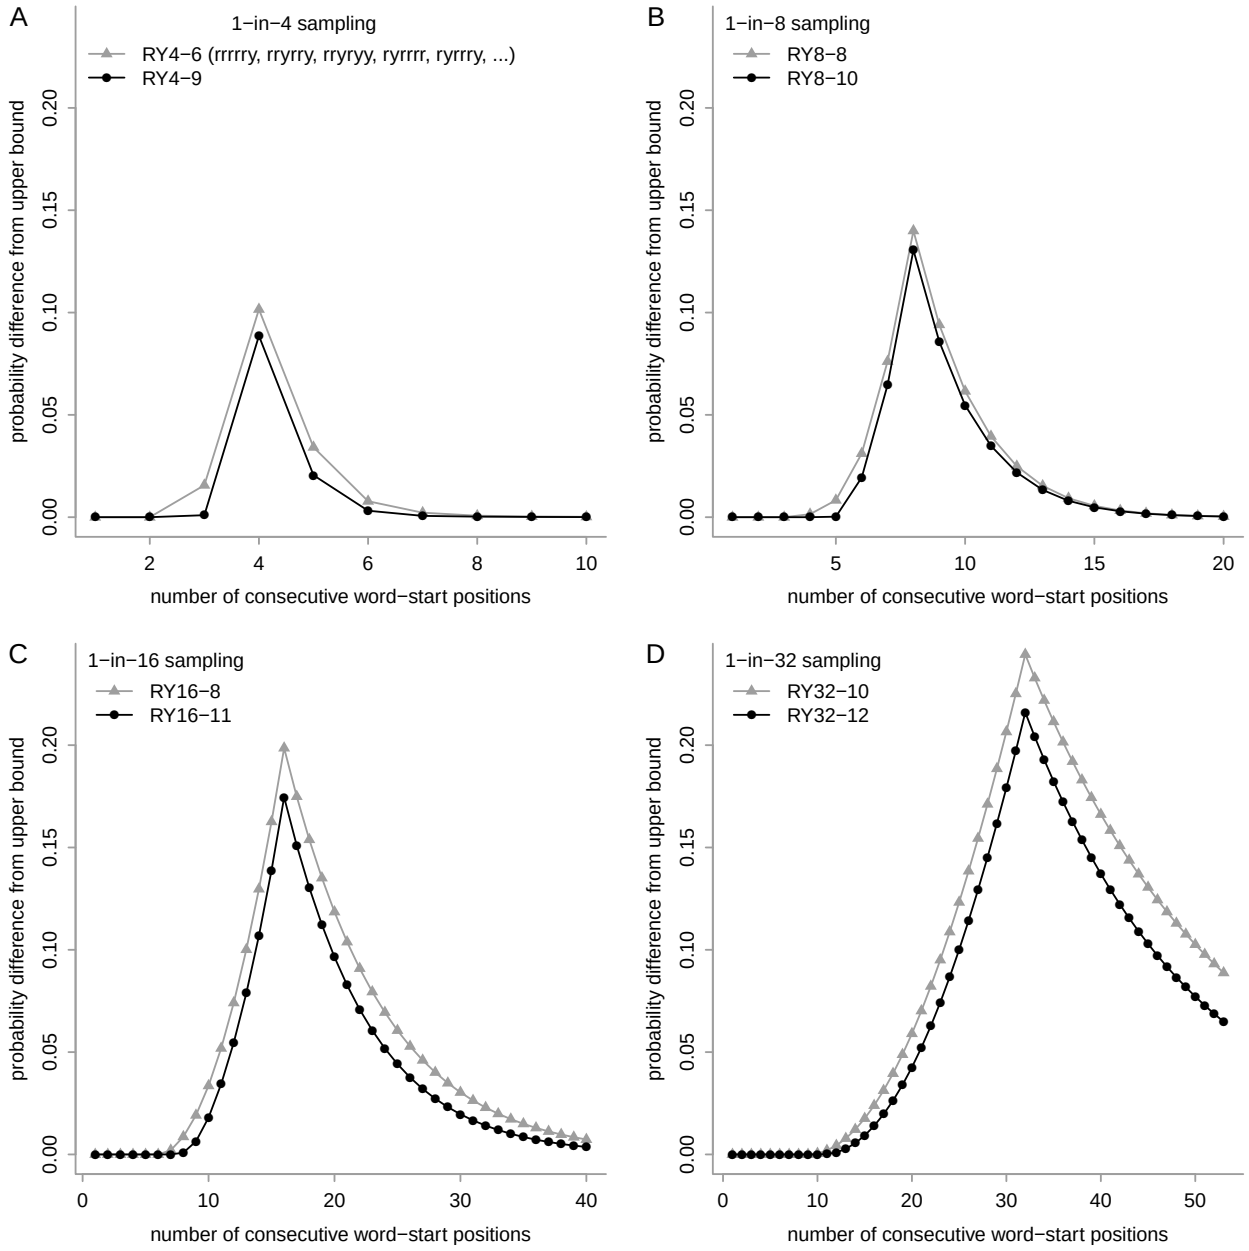

Figure 3: Run-hitting probabilities, for sampling schemes with sparsity: (**A**) 4, (**B**) 8, (**C**) 16, and (**D**) 32. For each run size (horizontal axis), the probability difference from the upper bound is shown (vertical axis). The gray triangles show the best schemes we found previously [11], and the black dots show new schemes. The triangles in panel **A** are equivalent to the triangles in fig. 2.

Table 1: Sets of `r`/`y` words analyzed in this study

name	sparsity	word length	number of words	min separation	max separation	count of `yr`	count of `ry`
old							
RY4-6	4	6	16	2	∞	17	25
RY8-8	8	8	32	3	∞	59	53
RY16-8	16	8	16	6	∞	28	31
RY32-10	32	10	32	9	∞	79	55
new							
RY4-9	4	9	128	2	∞	248	268
RY4push	4	9	128	3	∞	245	270
RY4pull	4	9	128	1	6	254	258
RY8-10	8	10	128	4	∞	268	314
RY16-11	16	11	128	7	∞	320	321
RY32-12	32	12	128	10	∞	309	384

Optimizing a weighted average

By optimizing the average run-hitting probability, we attach equal importance to each run length, which is not necessarily correct. Actually, short runs are more important for comparing highly-diverged sequences, and long runs are more important for comparing highly-similar sequences. This is because runs of overlapping matches tend to be shorter in more-diverged sequences.

So we also tried optimizing a weighted average of the run-hitting probabilities, where the weight decreases by a factor d for each increment in run length. For sparsity 4, with $d = 4$ we found a set of length-9 words named `RY4push` (fig. 4). These words exactly achieve the upper bound run-hitting probability for run size ≤ 3. This means that these words cannot occur at more than 1 out of 3 consecutive positions, in other words, they sample positions with minimum separation 3. This is optimal for sparsity 4: a minimum separation of 4 would require that the spacing is always exactly 4, which is impossible in e.g. a homopolymer sequence like `rrrrrrrrrr`. A word set with guaranteed minimum separation has been termed a "polar set" [30].

We also tried attaching higher weight to longer runs. For sparsity 4, with $d = 1/8$ we found a set of length-9 words named `RY4pull` (fig. 4). These words achieve the upper bound run-hitting probability, which is 1, for run size ≥ 6. In other words, they sample positions with a maximum separation of 6. Since these words are guaranteed to occur in a finite sequence, they are an example of a universal hitting set [22].

In practice, `RY4-9` is likely near-optimal for both closely- and distantly-related sequences, because its run-hitting probabilities are never much lower than those of `RY4push` and `RY4pull` (fig. 4).

One-hit specificity

Our approach increases the sensitivity of one-hit-based sequence comparison, but it's important to consider specificity too. For example, we can trivially increase sensitivity by using shorter matches, but that produces more matches between unrelated sequences, i.e. decreases specificity.

Is there any difference in specificity of `ry`- and `rr`-sampling? One way to measure specificity is by the number of sampled matches between unrelated random sequences: then there is no difference between `ry` and `rr`. However, one-hit seed-and-extend methods do not really attempt extension from every match: they avoid redundant extensions from overlapping matches. This suggests that two levels of specificity should be considered: all matches (which must be checked for

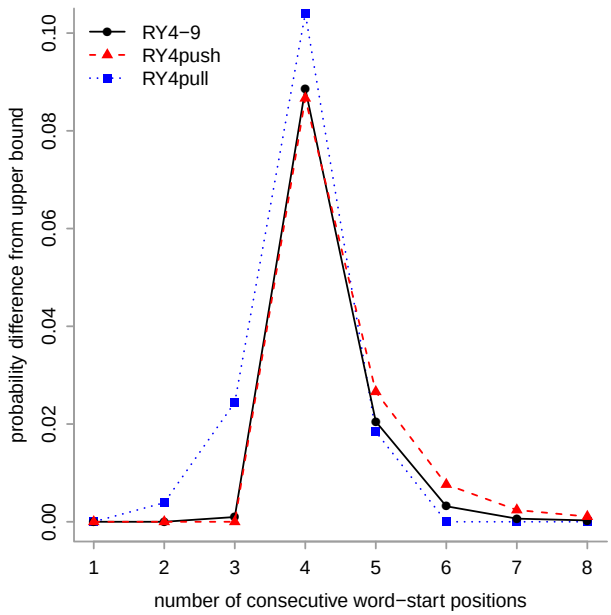

Figure 4: Run-hitting probabilities, for three 1-in-4 sampling schemes. For each run size (horizontal axis), the probability difference from the upper bound is shown (vertical axis). The black dots are the same as in fig. 3A.

overlap), and non-redundant matches (which trigger extension).

Let us investigate the number of sampled non-redundant matches between unrelated random sequences. The definition of non-redundant varies, but suppose that each match (of length m) is immediately extended into a maximally-long exact match, and subsequent matches in such extensions are skipped. Then, specificity can be measured by the number of sampled maximal exact matches between unrelated random sequences. The frequency of length-y maximal exact matches is proportional to p^y (where $p = ¼$ for DNA), and the probability of getting sampled is H_{y-m+1}. So the expected fraction of maximal exact matches (of length $\geq m$) that get sampled is

$$(1-p)\,\textstyle\sum_{x=1}^{\infty} H_x\, p^{x-1}\,. \qquad (4)$$

For example, `ry` samples a fraction 0.327 of maximal exact matches between unrelated random sequences, and `rr` samples 0.291. The upper bound, which applies to every 4-th sampling, is 0.332. So our approach has a downside: it maximizes non-redundant matches between unrelated sequences.

Methods other than one-hit

Our sequence-sampling approach is beneficial not only for one-hit-based sequence comparison. Firstly, we conjecture that it maximizes the sensitivity of two-hit methods. Their sensitivity depends on finding two nearby, non-overlapping seeds: these seeds could be in the same run, or in separate runs (fig. 1). Our approach optimizes the probability of hitting separate runs, and of non-overlapping hits in one run (see the Supplement). Thus, it seems likely to improve the sensitivity of two-hit methods.

Our approach is also likely to increase the sensitivity of coverage-based sequence comparison. That is because it provably increases the expected number of letters covered by matches in a true alignment. This is proportional to the probability that one letter is covered, which was shown [26] to equal

$$\textstyle\sum_{x=1}^{m} H_x\, P_{\text{run}}(x)\,, \qquad (5)$$

where m is the match length, and $P_{\text{run}}(x)$ is the probability that exactly x matches overlap one letter. Thus, higher run-hitting probabilities (H_x) imply higher expected coverage, for any level of sequence divergence.

Alignment-free methods are diverse, but a fundamental example is inferring the evolutionary distance between two sequences from the number of seed matches between them [18, 5]. To make this inference tight, we wish the variance in number of seed matches, for any given evolutionary distance, to be as low as possible [18]. Our sampling approach minimizes the variance of distance between sampled positions, which seems likely to reduce variance of match count.

Comparison to syncmers

A recent sequence-sampling approach is syncmers [9]: how does it compare? Syncmers are a special case of word-based sampling. To define whether a length-k word is a syncmer, we first map each of its length-j subwords to an integer, by an arbitrary hash function.

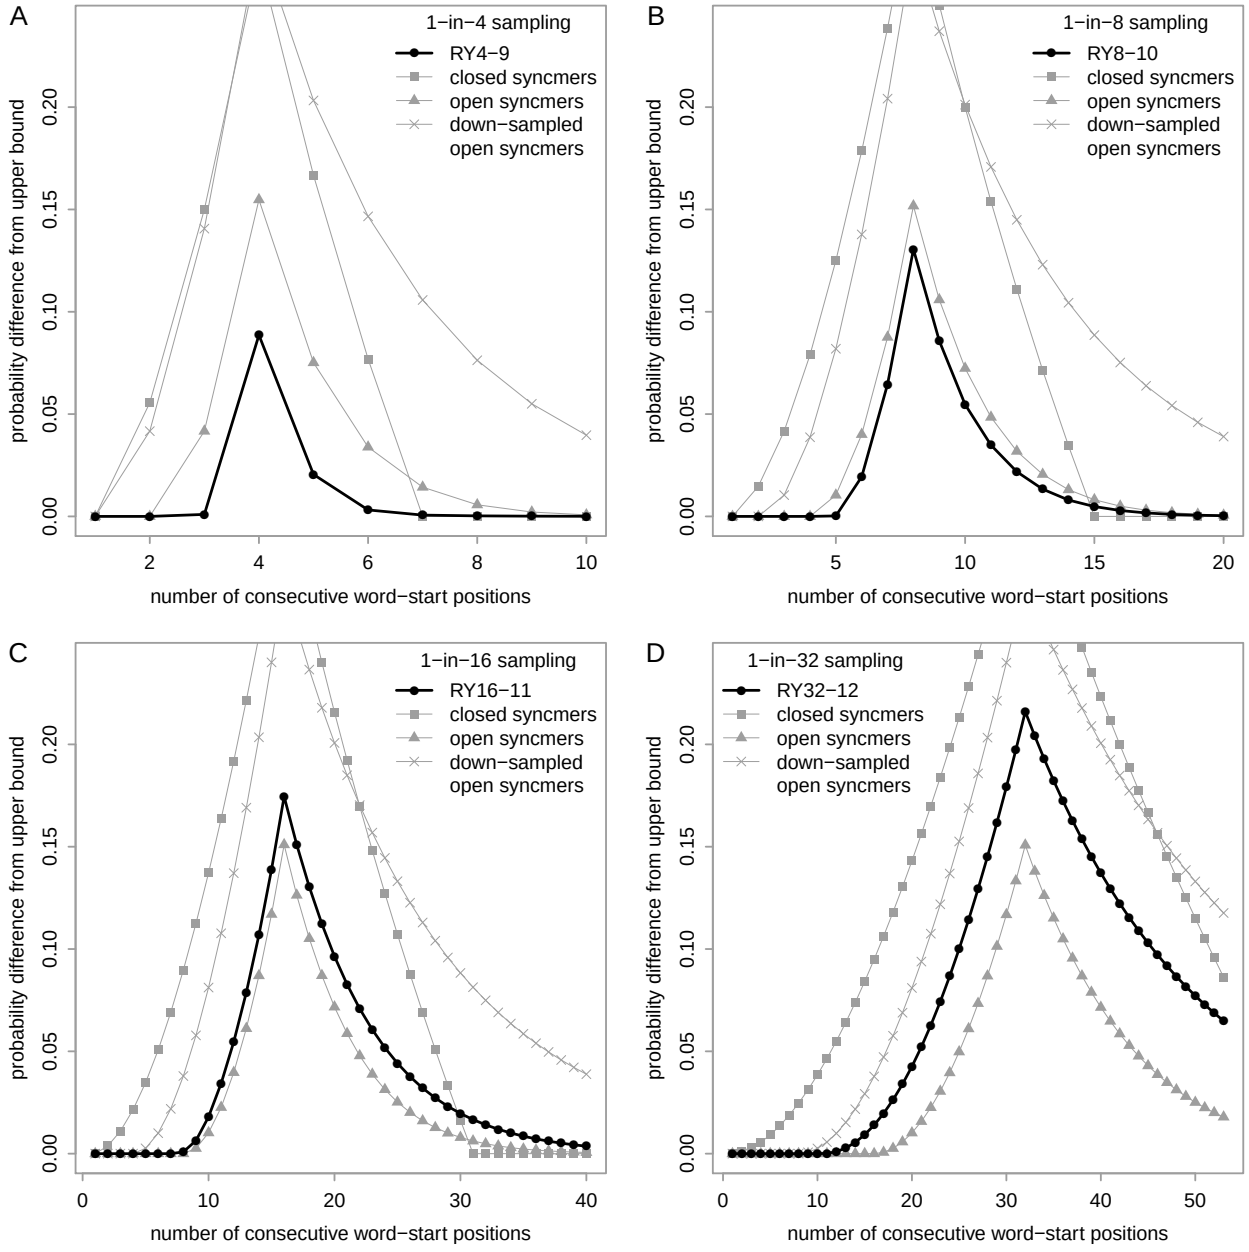

Figure 5: Run-hitting probabilities (difference from upper bound) for syncmers versus our word sets. The black dots are the same as in fig. 3. The down-sampled syncmers are 2-fold down-sampled.

The number of subwords is $w = k - j + 1$. If the minimum j-mer is the first or last one in the k-mer, then the k-mer is a "closed syncmer". If the minimum j-mer is the one at offset t ($1 \leq t \leq w$), then the k-mer is an "open syncmer with offset t". Finally, a d-fold "down-sampled" open or closed syncmer is one where the hash function of the k-mer is $\leq h/d$, where h is the maximum possible hash.

To make syncmer properties simple, we shall assume the subword size j is large enough that tied hash values are rare. Then, open syncmers sample 1 in w positions, and closed syncmers sample 2 in w positions. Also, the run-hitting probabilities of open syncmers are maximized when $t = \lceil w/2 \rceil$ [26], so we shall assume this value for t.

The run-hitting probabilities of syncmers are shown in fig. 5. Closed syncmers are universal hitting sets: for sampling sparsity s, they guarantee to hit runs of size $\geq 2s - 1$. For shorter runs, however, their run-hitting probabilities are inferior. Open syncmers are polar word sets: they have minimum separation $\lceil s/2 \rceil$. They are less sensitive than our word sets for sparsity 4 and 8, but more sensitive for sparsity 16 and 32. It is not too surprising that these open syncmers are more sensitive than our word sets, because they use longer words. The word length of open syncmers, for sparsity s, is $> s$. Thus, for sparsity 16 or 32, open syncmers can only be used with rather long matches, which have low sensitivity. Down-sampled open syncmers can use shorter words, but their run-hitting probabilities are inferior (fig. 5). Overall, open syncmers are a good heuristic for maximizing run-hitting probabilities.

Sparsity in biological DNA

To see if our sampling methods are likely to work well for real DNA, we first checked the sparsity of our word sets in biological data. We checked human DNA (the first million non-`N` bases of hg38 chromosome 2), AT-rich *Plasmodium falciparum* DNA (the first million bases of NC_000521.4), and repetitive centromeric DNA (the 161 kb sequence 000736F from medaka strain Hd-rR 2.2.4 [12]). For the human and centromeric DNA, the words' sparsity is close to the expected values of 4, 8, 16, and 32 (fig. 6). For the plasmodium DNA, some word sets are sparser than expected, but not wildly different. This suggests that our sampling schemes are applicable to a wide range of natural DNA.

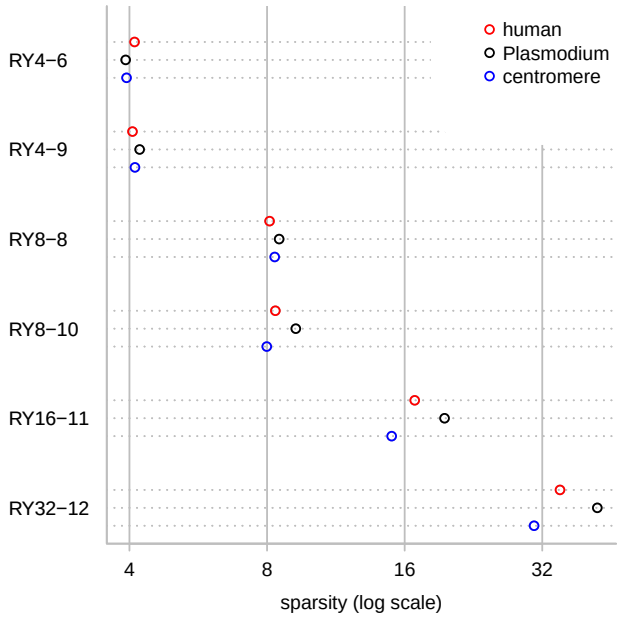

Figure 6: Sparsity of some sampling schemes in three biological DNA sequences.

Specificity in biological DNA

We then measured specificity in real DNA, by counting the number of sampled matches between two unrelated sequences: 10^6 bases of human DNA (the same as above) and 10^6 bases of reversed (but not complemented) mouse DNA (the first million non-`N` bases of mm39 chromosome 1). Some sampling schemes find more matches than others (fig. 7), so they are less specific. For example, `RY4-6` finds more matches than `RY4-9`: this is because `RY4-6` includes the word `ryryry`, and that word alone gets most of the long matches found by `RY4-6`. This is presumably because `ryryry` samples simple repeats like `atatat`, which are frequent in real DNA. Similarly, `RY8-8` finds more matches than `RY8-10`, because `RY8-8` includes the repetitive words `yyyryyyr` and `yyryyryy`. By avoiding simple repeats, our new word sets are more specific

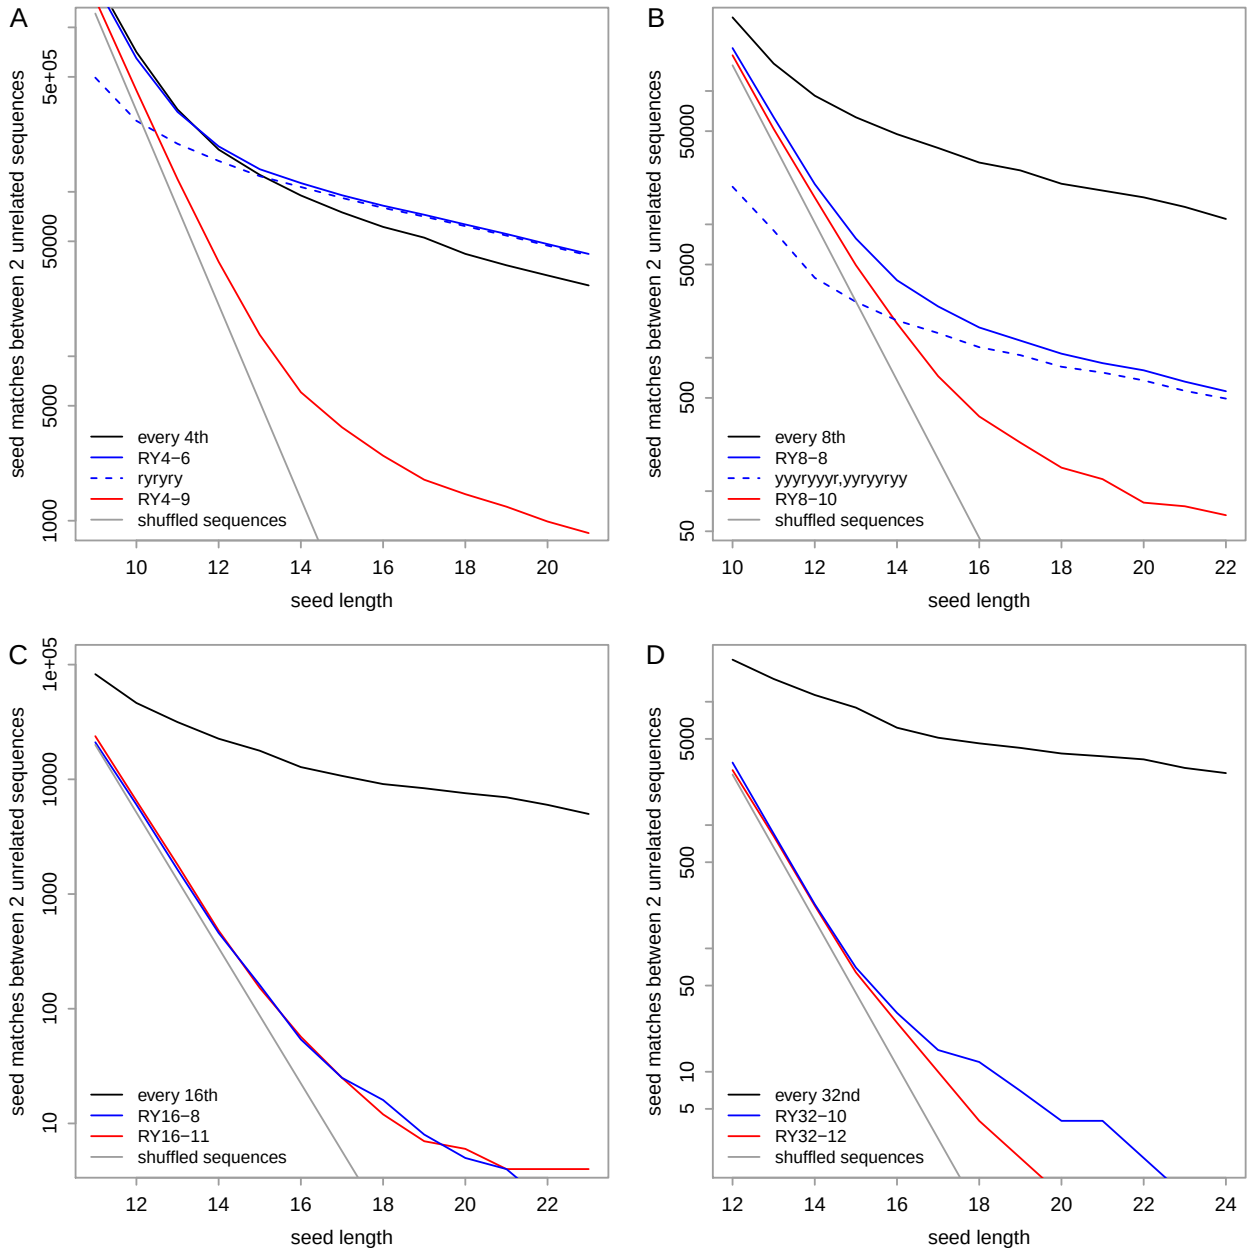

Figure 7: Specificity of some seed-sampling schemes, for mammal DNA. The vertical axis is the number of pair-wise seed matches between: 10^6 bases of human DNA and 10^6 bases of reversed (but not complemented) mouse DNA. In each panel, the gray line shows the expected number of hits if each sequence were shuffled.

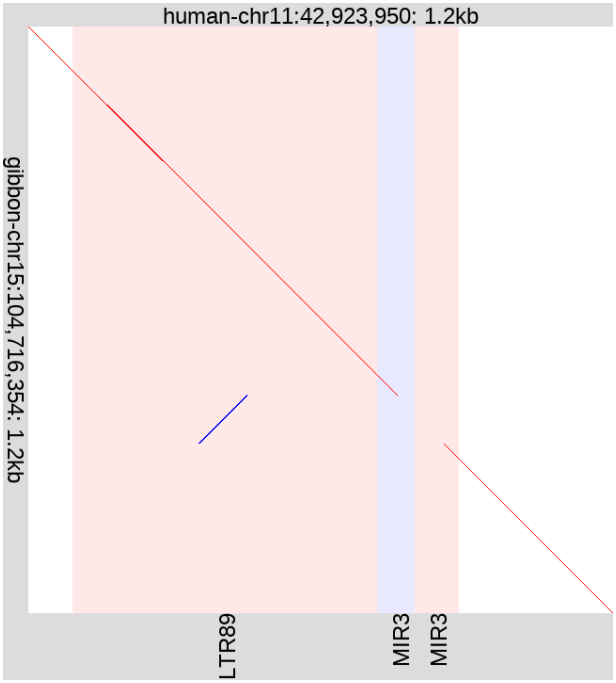

Figure 8: A short-range rearrangement in the gibbon genome relative to human. The diagonal lines show alignments: red = same strand, blue = opposite strand. The vertical stripes show repeat elements in the human sequence: red = forward-oriented, blue = reverse-oriented. By visually scanning this figure from top to bottom, we can see which part of the human sequence each part of the gibbon sequence comes from.

than theoretically-optimal every s-th sampling.

Sensitivity in biological DNA

Finally, we did one test of sensitivity in real DNA, by seeking short-range rearrangements in the gibbon genome relative to the human genome (e.g. fig. 8). Specifically, we sought rearrangements that may arise by short-range template switching during DNA replication, which are often overlooked [14].

We used a two-step genome alignment method with the LAST software [10]. First, it finds and aligns similar segments of the two genomes. Second, it cuts these alignments down to a unique best alignment for each part of the gibbon genome. This is a reasonable way to find rearrangements that occurred in the gibbon lineage. The first step uses one-hit seed-and-extend, with variable-length seeds: at each sampled position in the gibbon sequence, it uses the shortest seed that occurs at most m times in human. Here we used exact-match seeds with $m = 2$.

As our ground truth, we used gibbon-human alignments found by the same method without any sampling. This will not be perfectly accurate, but it suffices to compare sampling methods' abilities to find related segments. We defined "short range" rearrangements by arbitrary thresholds (see the Supplement), and deemed a sampling method to find a rearrangement if it finds one with the same start and end coordinates of the rearranged fragment.

In this test, the sensitivities of the sampling schemes agree with theoretical expectations. Firstly, every s-th sampling always has the highest sensitivity, e.g. every-16th sampling misses fewer rearrangements than `RY16-11` (fig. 9A). Secondly, our new word sets are always more sensitive than our old word sets (table 1), e.g. `RY4-9` misses fewer rearrangements than `RY4-6`. Furthermore, `RY4-6` is more sensitive than `RY`. Finally, the word sets are more sensitive when oriented to minimize `yr` count, in most cases.

The run times show the benefit of more-aggressive sampling (fig. 9B). Every s-th seeding used more memory, because it uses an index data-structure of all positions in the human genome, whereas word-based sampling just indexes the sampled positions. Among the 1-in-4 sampling schemes, RY and YR are fastest: we speculate this is because they have lower sensitivity so find fewer alignments.

Discussion

We have presented an approach to sampling a sequence, which provably optimizes sensitivity for one-hit-based sequence comparison, for idealized random DNA. It is likely near-optimal for a wide range of alignment-based and alignment-free methods. For real biological DNA, our approach has close to the expected sampling rates, improves sensitivity in a limited test, and even improves specificity by avoiding simple repeats. Thus, our approach is promising for

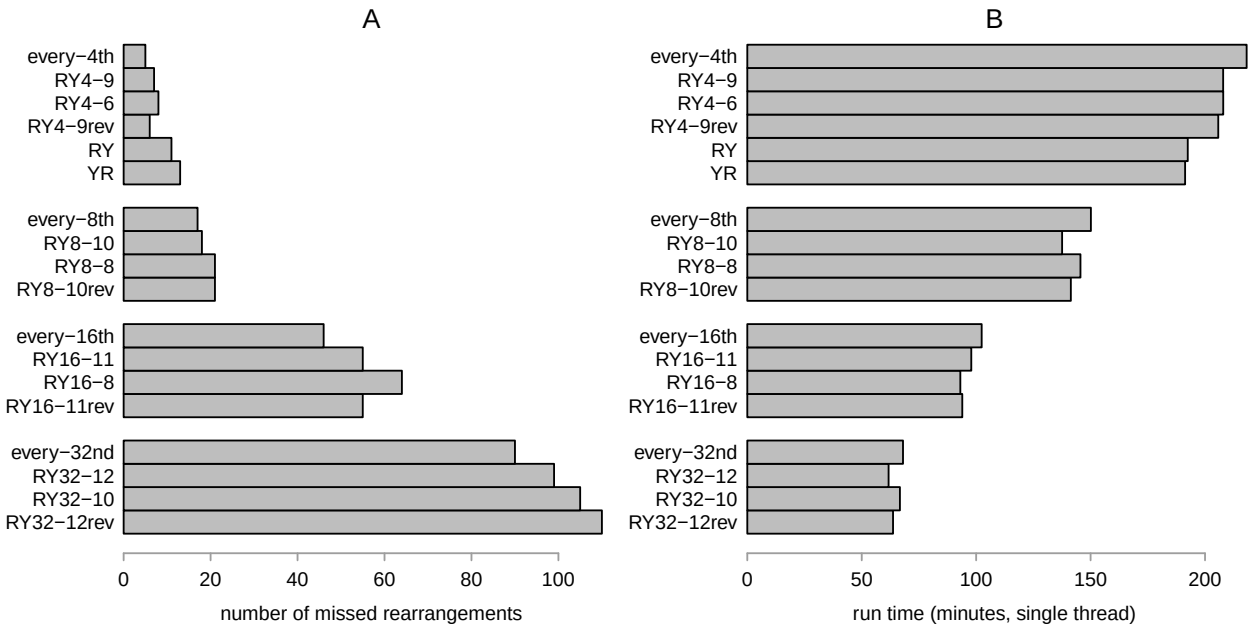

Figure 9: Sensitivity for finding short-range rearrangements, and run time, when aligning human and gibbon genomes using LAST with various seed-sampling schemes. The suffix rev indicates sets of reversed words, which thus maximize `yr` count.

real DNA.

Our sequence-sampling approach is a balance between universal hitting sets and polar sets. Universal hitting sets guarantee to sample a sufficiently-long sequence at least once, whereas polar sets have optimal hitting probabilities for sufficiently-short sequences which they sample at most once. It is impossible, however, for a word set to be both universal and polar. This is because a polar set with minimum separation b must exclude repeats (like `ryryry`) with period $< b$, but a universal set must include them. So it may be interesting to find "near-universal hitting sets" that hit every sequence except short-period repeats: such a set could perhaps be polar, with the specificity benefit of avoiding simple repeats.

There are several other directions for future advances. Firstly, it would be useful to find better-optimized word sets than we found by simulated annealing, with longer words. Related to this, it would be interesting to know if optimal or near-optimal word sets have any simple properties. It would be useful to optimize sampling of inexact seeds [11], which we feel are under-used, though they benefit alignment-based and alignment-free methods (e.g. [15, 18, 23]). In particular, transition seeds, which tolerate frequently-occurring `a`↔`g` and `c`↔`t` substitutions, are greatly neglected and promising for improved sequence comparison [20]. It would also be useful to optimize protein sequence sampling, perhaps using a reduced alphabet. Finally, it may be useful to optimize sampling of specific sequences, such as a reference human genome [6, 30]: it will be interesting to see how much this improves over sampling methods that are optimal for random sequences.

We have clarified how to do optimal sampling of a sequence, which is a fundamental way to rapidly analyze big sequence data. At the same time, we have shown that random sequences of independent letters are surprisingly interesting.

Acknowledgements

We thank Anish Shrestha and Kiyoshi Asai for helpful comments, and Animesh Awasthi for a comparison of sequence-sampling methods which motivated this work.

References

[1] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman. Basic local alignment search tool. *Journal of molecular biology*, 215(3):403–410, 1990.

[2] S. F. Altschul, T. L. Madden, A. A. Schäffer, J. Zhang, Z. Zhang, W. Miller, and D. J. Lipman. Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic acids research*, 25(17):3389–3402, 1997.

[3] G. Benson and D. Y. Mak. Exact distribution of a spaced seed statistic for DNA homology detection. In *International Symposium on String Processing and Information Retrieval*, pages 282–293. Springer, 2008.

[4] J. Bérard, J.-B. Gouéré, and D. Piau. Solvable models of neighbor-dependent substitution processes. *Mathematical Biosciences*, 211(1):56–88, 2008.

[5] G. Bernard, C. X. Chan, Y.-b. Chan, X.-Y. Chua, Y. Cong, J. M. Hogan, S. R. Maetschke, and M. A. Ragan. Alignment-free inference of hierarchical and reticulate phylogenomic relationships. *Briefings in Bioinformatics*, 20(2):426–435, 2019.

[6] R. Chikhi, A. Limasset, S. Jackman, J. T. Simpson, and P. Medvedev. On the representation of de Bruijn graphs. In *International conference on Research in computational molecular biology*, pages 35–55. Springer, 2014.

[7] M. Csűrös. Performing local similarity searches with variable length seeds. In *Annual Symposium on Combinatorial Pattern Matching*, pages 373–387. Springer, 2004.

[8] A. Dutta, D. Pellow, and R. Shamir. Parameterized syncmer schemes improve long-read mapping. *bioRxiv*, 2022.

[9] R. Edgar. Syncmers are more sensitive than minimizers for selecting conserved k-mers in biological sequences. *PeerJ*, 9:e10805, 2021.

[10] M. C. Frith and R. Kawaguchi. Split-alignment of genomes finds orthologies more accurately. *Genome biology*, 16(1):1–17, 2015.

[11] M. C. Frith, L. Noé, and G. Kucherov. Minimally overlapping words for sequence similarity search. *Bioinformatics*, 36(22-23):5344–5350, 2020.

[12] K. Ichikawa, S. Tomioka, Y. Suzuki, R. Nakamura, J. Yoshimura, M. Kumagai, Y. Inoue, Y. Uchida, N. Irie, H. Takeda, et al. Centromere evolution and CpG methylation during vertebrate speciation. *Nature communications*, 8(1):1–11, 2017.

[13] B. Linard, K. Swenson, and F. Pardi. Rapid alignment-free phylogenetic identification of metagenomic sequences. *Bioinformatics*, 35(18):3303–3312, 2019.

[14] A. Löytynoja and N. Goldman. Short template switch events explain mutation clusters in the human genome. *Genome research*, 27(6):1039–1049, 2017.

[15] B. Ma, J. Tromp, and M. Li. PatternHunter: faster and more sensitive homology search. *Bioinformatics*, 18(3):440–445, 2002.

[16] U. Manber. Finding similar files in a large file system. In *USENIX winter 1994 technical conference*, volume 94, pages 1–10, 1994.

[17] B. Morgenstern. Sequence comparison without alignment: The SpaM approaches. In *Multiple Sequence Alignment*, pages 121–134. Springer, 2021.

[18] B. Morgenstern, B. Zhu, S. Horwege, and C. A. Leimeister. Estimating evolutionary distances between genomic sequences from spaced-word matches. *Algorithms for Molecular Biology*, 10(1):1–12, 2015.

[19] G. Myers. Efficient local alignment discovery amongst noisy long reads. In *International Workshop on Algorithms in Bioinformatics*, pages 52–67. Springer, 2014.

[20] L. Noé and G. Kucherov. Improved hit criteria for DNA local alignment. *BMC bioinformatics*, 5(1):1–9, 2004.

[21] L. Noé and D. E. Martin. 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.

[22] Y. Orenstein, D. Pellow, G. Marçais, R. Shamir, and C. Kingsford. Compact universal k-mer hitting sets. In *International Workshop on Algorithms in Bioinformatics*, pages 257–268. Springer, 2016.

[23] R. Ounit and S. Lonardi. Higher classification sensitivity of short metagenomic reads with CLARK-S. *Bioinformatics*, 32(24):3823–3825, 2016.

[24] M. Roberts, W. Hayes, B. R. Hunt, S. M. Mount, and J. A. Yorke. Reducing storage requirements for biological sequence comparison. *Bioinformatics*, 20(18):3363–3369, 2004.

[25] S. Schleimer, D. S. Wilkerson, and A. Aiken. Winnowing: local algorithms for document fingerprinting. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*, pages 76–85, 2003.

[26] J. Shaw and Y. W. Yu. Theory of local k-mer selection with applications to long-read alignment. *Bioinformatics*, 2021.

[27] S. Vinga and J. Almeida. Alignment-free sequence comparison—a review. *Bioinformatics*, 19(4):513–523, 2003.

[28] D. E. Wood, J. Lu, and B. Langmead. Improved metagenomic analysis with Kraken 2. *Genome biology*, 20(1):1–13, 2019.

[29] H. Yi and L. Jin. Co-phylog: an assembly-free phylogenomic approach for closely related organisms. *Nucleic acids research*, 41(7):e75–e75, 2013.

[30] H. Zheng, C. Kingsford, and G. Marçais. Sequence-specific minimizers via polar sets. *Bioinformatics*, 37(Supplement_1):i187–i195, 2021.

[31] W. Zhou, G. Liang, P. L. Molloy, and P. A. Jones. DNA methylation enables transposable element-driven genome expansion. *Proceedings of the National Academy of Sciences*, 117(32):19359–19366, 2020.